



US009058695B2

(12) **United States Patent**
Cardno et al.

(10) **Patent No.:** **US 9,058,695 B2**
(45) **Date of Patent:** **Jun. 16, 2015**

(54) **METHOD OF GRAPHICALLY
REPRESENTING A TREE STRUCTURE**

(58) **Field of Classification Search**

None

See application file for complete search history.

(75) Inventors: **Andrew John Cardno**, San Diego, CA
(US); **Peter Stewart Ingham**, Avalon
(NZ); **Bart Andrew Lewin**, Woodland
Hills, CA (US); **Ashok Kumar Singh**,
Henderson, NV (US)

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,694,594 A * 12/1997 Chang 1/1
5,708,767 A * 1/1998 Yeo et al. 345/440

(Continued)

(73) Assignee: **New BIS Safe Luxco S.A R.L.**,
Luxembourg (LU)

(*) Notice: Subject to any disclaimer, the term of this
patent is extended or adjusted under 35
U.S.C. 154(b) by 690 days.

FOREIGN PATENT DOCUMENTS

WO WO 2004/068300 8/2004
WO WO 2006/138502 12/2006

(21) Appl. No.: **13/000,315**

OTHER PUBLICATIONS

(22) PCT Filed: **Jun. 18, 2009**

Barbec et al. "Visualizing and Animating R-trees and Spatial Opera-
tions in Spatial Databases on the Worldwide Web." *Visual Database
Systems 4 (VBD4)*. 1998. pp. 147-153.

(86) PCT No.: **PCT/NZ2009/000110**

§ 371 (c)(1),
(2), (4) Date: **Mar. 29, 2011**

(Continued)

(87) PCT Pub. No.: **WO2009/154480**

Primary Examiner — Stephen Alvesteffier

(74) *Attorney, Agent, or Firm* — Merchant & Gould P.C.

PCT Pub. Date: **Dec. 23, 2009**

(57)

ABSTRACT

(65) **Prior Publication Data**

US 2011/0179370 A1 Jul. 21, 2011

Related U.S. Application Data

(60) Provisional application No. 61/074,347, filed on Jun.
20, 2008, provisional application No. 61/115,036,
filed on Nov. 15, 2008, provisional application No.
61/140,556, filed on Dec. 23, 2008.

(51) **Int. Cl.**

G06F 3/048 (2013.01)

G06T 11/20 (2006.01)


G06F 17/30 (2006.01)

(52) **U.S. Cl.**

CPC **G06T 11/206** (2013.01); **G06F 17/30994**
(2013.01)

In a data visualization system, a method of creating a visual
representation of data points from metric data, wherein the
method includes the steps of: a data retrieval module retriev-
ing the metric data from a data storage module in communica-
tion with the data visualization system, a determination
module determining the data points based on the retrieved
metric data, and a display module arranging the data points
for displaying on a display device according to a predeter-
mined visual representation, arranging the data points into a
plurality of meta groups in a hierarchical manner, and arrang-
ing the meta groups into a plurality of layers, where each layer
represents the data points at different levels of granularity,
and a statistical distance determination module determining a
statistical distance between items in the meta groups, and the
display module arranging the items within the meta groups
based on the determined statistical distance.

35 Claims, 15 Drawing Sheets



(56)

References Cited**U.S. PATENT DOCUMENTS**

6,347,318	B1	2/2002	Rokicki	
6,463,400	B1	10/2002	Barkley-Yeung	
6,499,026	B1 *	12/2002	Rivette et al.	1/1
6,583,794	B1 *	6/2003	Wattenberg	715/708
6,765,597	B2 *	7/2004	Barksdale et al.	715/853
6,888,543	B2 *	5/2005	Ingber et al.	345/420
7,287,033	B2 *	10/2007	Shadmon et al.	1/1
7,333,998	B2	2/2008	Heckerman et al.	
2002/0193981	A1	12/2002	Keung et al.	
2003/0200191	A1	10/2003	Pao et al.	
2004/0181554	A1	9/2004	Heckerman et al.	
2005/0060329	A1	3/2005	Watanabe et al.	
2007/0097965	A1	5/2007	Qiao et al.	
2007/0124670	A1 *	5/2007	Finck et al.	715/526
2008/0059872	A1 *	3/2008	Tseng et al.	715/231
2008/0071843	A1	3/2008	Papadimitriou et al.	
2008/0104532	A1 *	5/2008	Stambaugh	715/771
2008/0134029	A1 *	6/2008	Shelton et al.	715/700
2008/0140644	A1 *	6/2008	Franks et al.	707/5
2009/0287991	A1 *	11/2009	Nakamura et al.	715/210

OTHER PUBLICATIONS

Brabec et al. "The VASCO R-tree JAVA Applet." *Visual Database Systems 4 (VBD4)*. 1998. pp. 123-140.

Guttman. "R-Trees: A Dynamic Index Structure for Spatial Searching." *ACM SIGMOD International Conference on Management of Data*. 1984. pp. 47-57.

Wikipedia; "Heat map;" Date Accessed: Jun. 10, 2008; http://en.wikipedia.org/wiki/Heat_map.

Nasdaq; "Nasdaq-100 Dynamic Heatmap;" Date Accessed: Jun. 10, 2008; http://screening.nasdaq.com/heatmaps/heatmap_100.asp.

Weather.com; Date Accessed: Jan. 31, 2008; <http://weather.com/>.

Information Aesthetics; "travel time maps;" Date Accessed: Jan. 31, 2008; <http://infosthetics.com/archives/locative/>.

Few, Stephen; from white paper "BizViz: The Power of Visual Business Intelligence;" Mar. 7, 2006; www.perceptualedge.com.

Wikipedia; "Voronoi Diagram;" Date Accessed: Jan. 31, 2008; http://en.wikipedia.org/wiki/Voronoi_diagram.

Entropia; "Más tiempo;" Date Accessed: Jan. 31, 2008; <http://www.luispabon.com/entropia/index.php?entry=entry071129-145959>.

Wikipedia; "Agile Software Development;" Date Accessed: Jan. 30, 2008; http://en.wikipedia.org/wiki/Agile_software_development.

Wikipedia; "Dimensional Modeling;" Date Accessed: Apr. 10, 2008; http://en.wikipedia.org/wiki/Dimensional_modeling.

Kimball, Ralph; "A Dimensional Modeling Manifesto;" Date Accessed: Apr. 10, 2008; <http://www.dbmsmag.com/9708d15.html>.

Sam's Publishing; developer.com Gamelan™; "Introduction to EJB's;" <http://www.developer.com/java/ejb/article.php/1434371>, Jul. 29, 2002.

Gilbert, Cheryl, et. al.; SearchCIO—Midmarket; "IIOP;" Date Accessed: Jan. 30, 2008; http://searchcio-midmarket.techtarget.com/sDefinition/0,,sid183_gci214019,00.html.

Google; Google Maps; "KML Gallery: Explore the Earth on Google;" Date Accessed: Jan. 30, 2008; <http://earth.google.com/gallery/>.

Wikipedia; "Mean down time;" Date Accessed: Jan. 30, 2008; http://en.wikipedia.org/wiki/Mean_down_time.

Wikipedia; "Mean time between failures;" Date Accessed: Jan. 30, 2008; http://en.wikipedia.org/wiki/Mean_time_between_failures.

Wikipedia; "Mean time to recovery;" Date Accessed: Jan. 30, 2008; http://en.wikipedia.org/wiki/Mean_time_to_recovery.

Paris Technologies, Inc.; "OLAP;" Date Accessed: Jan. 30, 2008; <http://www.olap.com>.

Wikipedia; "Planogram;" Date Accessed: Jan. 30, 2008; <http://en.wikipedia.org/wiki/Planogram>.

Wikipedia; "Software as a Service;" Date Accessed: Jan. 30, 2008; http://en.wikipedia.org/wiki/Software_as_a_Service.

Wikipedia; "Self-organizing map;" Date Accessed: Jan. 30, 2008; http://en.wikipedia.org/wiki/Self-organizing_map.

Zeiger, Stefan; "Servlet Essentials," Version 1.3.6, Nov. 4, 1999; Date Accessed: Jan. 30, 2008; <http://www.novocode.com/doc/servlet-essentials/>.

Six Sigma; "Subject Matter Expert—SME;" Date Accessed: Jan. 30, 2008; http://www.isixsigma.com/dictionary/Subject_Matter_Expert_-_SME-396.htm.

IBM; WebSphere Product pp.; "WebSphere software;" Date Accessed: Jan. 30, 2008; http://www-306.ibm.com/software/websphere/?pgel=ibmhzn&cm_re=masthead_-_products_-_sw-websphere.

* cited by examiner

Nasdaq prices valid as of Jun. 9, 2008 Market Closed									
QQQQ									
-0.53%									
VRTX	STLD	RIMM	DELL	CTSH	ATVI	EXPD	HANS	SRCL	SYMC
2.54%	2.53%	2.08%	1.73%	1.73%	1.65%	1.62%	1.57%	1.33%	1.24%
BRCM	MCHP	NTAP	MSFT	LLTC	FWLT	JNPR	ADSK	JOYG	DISCA
1.21%	1.15%	1.05%	0.80%	0.77%	0.75%	0.75%	0.64%	0.55%	0.55%
YHOO	PAYX	SHLD	ADBE	ORCL	CHRW	CEPH	WFMI	AMAT	COST
0.53%	0.51%	0.46%	0.40%	0.40%	0.39%	0.36%	0.35%	0.21%	0.20%
AMGN	CELG	FLEX	CMCSA	FAST	LVLTL	SIRI	CDNS	DTV	MICC
0.14%	0.13%	0.09%	0.09%	0.06%	0.00%	0.00%	0.00%	0.00%	-0.02%
CA	QCOM	PCAR	SPLS	INTU	PDCO	CHKP	PETM	VMED	CTAS
-0.04%	-0.06%	-0.08%	-0.09%	-0.14%	-0.15%	-0.16%	-0.22%	-0.25%	-0.28%
XLNX	IACI	EXPE	MRVL	SIAL	CSCO	INTC	TEVA	LINTA	EBAY
-0.37%	-0.38%	-0.40%	-0.41%	-0.41%	-0.41%	-0.62%	-0.64%	-0.65%	-0.65%
DISH	GILD	SBUX	RYAAY	HOLX	VRSN	LBTYA	CTXS	ALTR	XRAY
-0.74%	-0.84%	-0.85%	-0.87%	-0.94%	-1.00%	-1.10%	-1.12%	-1.13%	-1.15%
INFY	FISV	GENZ	JAVA	ERTS	LRCX	MNST	AMZN	KLAC	NVDA
-1.18%	-1.28%	-1.30%	-1.30%	-1.30%	-1.34%	-1.39%	-1.49%	-1.52%	-1.54%
GOOG	NIHD	HSIC	AKAM	BIIB	ESRX	LEAP	ISRG	AAPL	BBBY
-1.61%	-1.64%	-1.69%	-1.79%	-1.83%	-1.90%	-2.06%	-2.07%	-2.17%	-2.26%
LOGI	APOL	GRMN	FMCN	BIDU	SNDK	WYNN	LAMP	UAUA	AMLN
-2.63%	-2.68%	-3.41%	-3.47%	-3.78%	-3.91%	-4.05%	-4.16%	-4.17%	-10.32%

Nasdaq prices valid as of Jun. 9, 2008 Market Closed																		
QQQQ																		
-0.53%																		
VRTX	STLD	RIMM	DELL	CTSH	ATVI	EXPD	HANS	SRCL	SYMC									
2.54%	2.53%	2.06%	1.73%	1.73%	1.65%	1.62%	1.57%	1.33%	1.24%									
BRCM	MCHP	Steel Dynamics, Inc. (STLD) Last Sale 38.97 Net (Percent) Change: 0.96 (2.53%) Today's High : 39.15 Today's Low: 38.01 Volume: 3,430,700 Previous Close: 38.01 52 Week Range: 16.81 - 39.56					ADSK	JOYG	DISCA									
1.21%	1.15%																	
YHOO	PAYX																	
0.53%	0.51%																	
AMGN	CELG																	
0.14%	0.13%																	
CA	QCOM																	
-0.04%	-0.06%																	
XLNX	IACI																	
-0.37%	-0.38%																	
DISH	GILD																	
-0.74%	-0.84%																	
INFY	FISV																	
-1.18%	-1.28%																	
GOOG	NIHD																	
-1.61%	-1.64%																	
LOGI	APOL																	
-2.63%	-2.68%																	
							WFMJ	AMAT	COST									
							0.35%	0.21%	0.20%									
							CDNS	DTV	MICC									
							0.00%	0.00%	-0.02%									
							PETM	VMED	CTAS									
							-0.22%	-0.25%	-0.28%									
							TEVA	LINTA	EBAY									
							-0.64%	-0.65%	-0.65%									
							CTXS	ALTR	XRAY									
							-1.12%	-1.13%	-1.15%									
							AMZN	KLAC	NVDA									
							-1.48%	-1.52%	-1.54%									
							ISRG	AAPL	BBBY									
							-2.07%	-2.17%	-2.26%									
							LAMP	UAUA	AMLN									
							-4.16%	-4.17%	-10.32%									
							</											

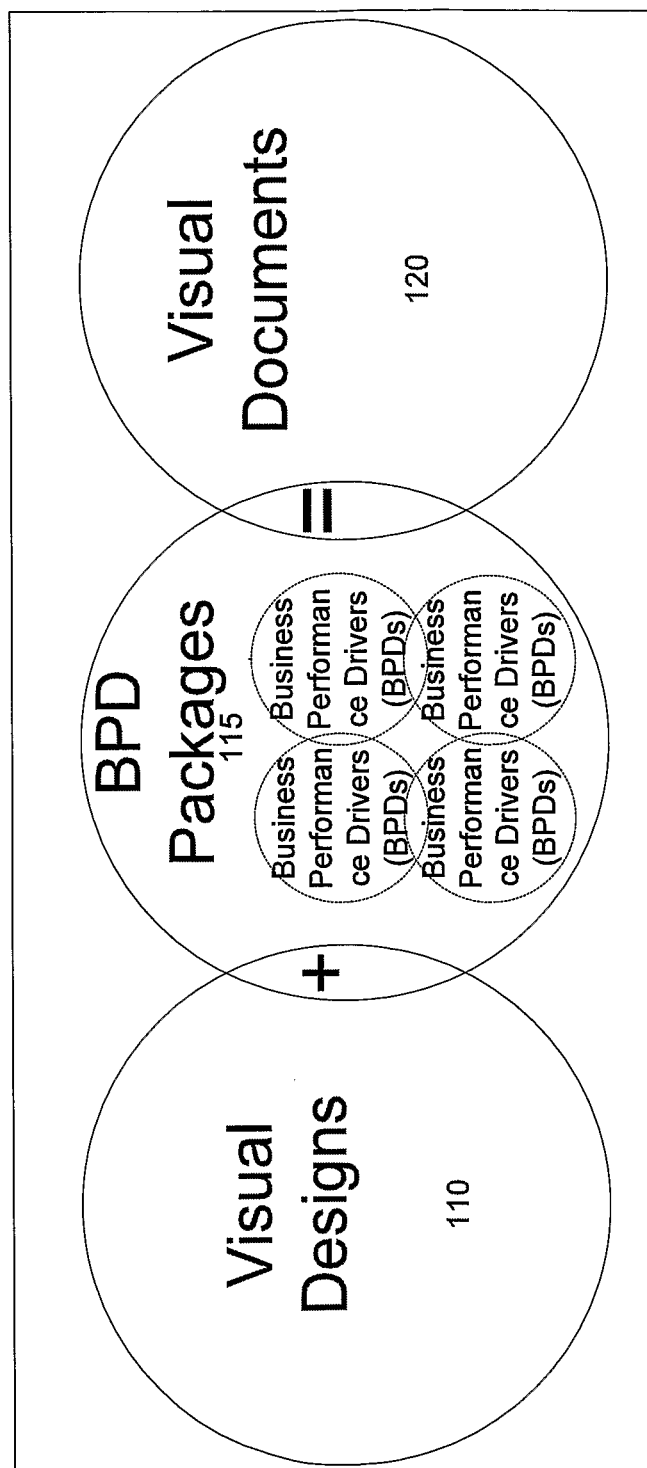


Figure 1C

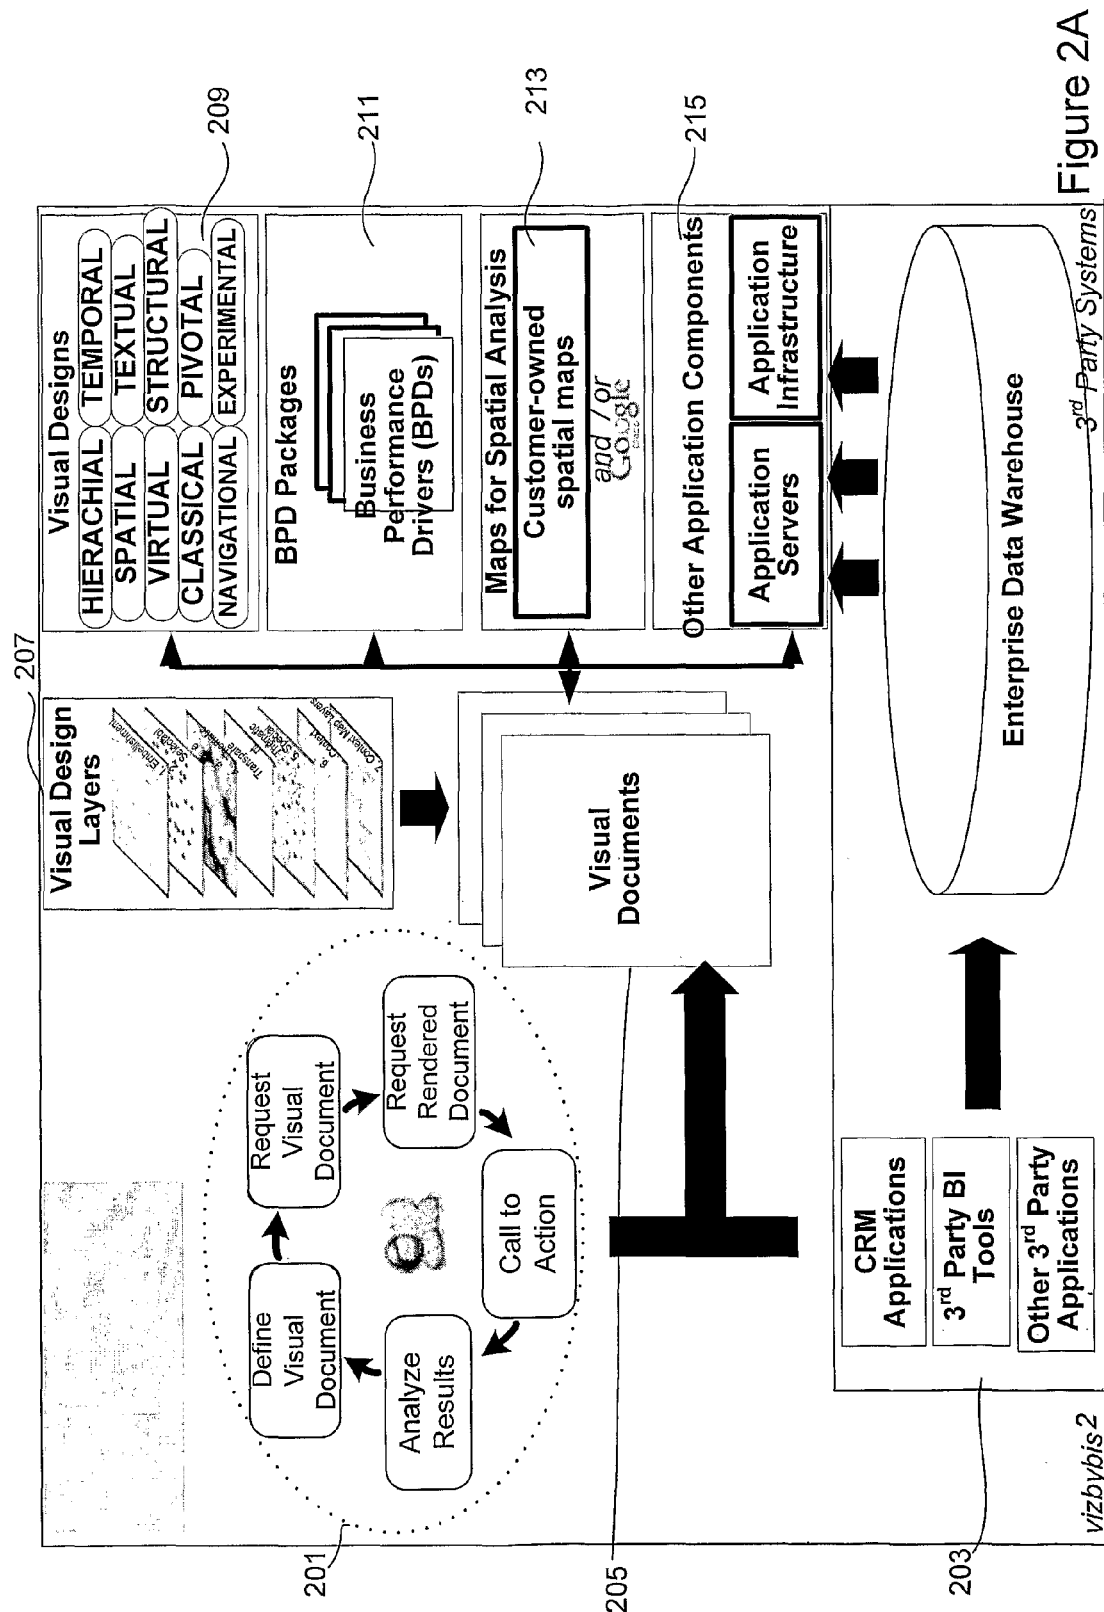


Figure 2A

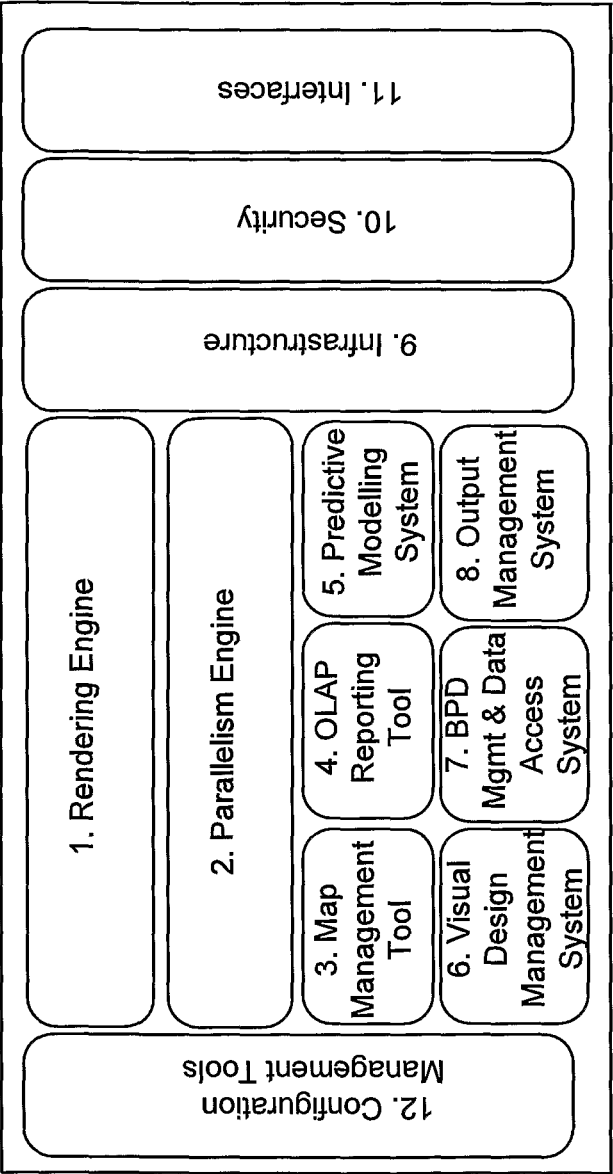
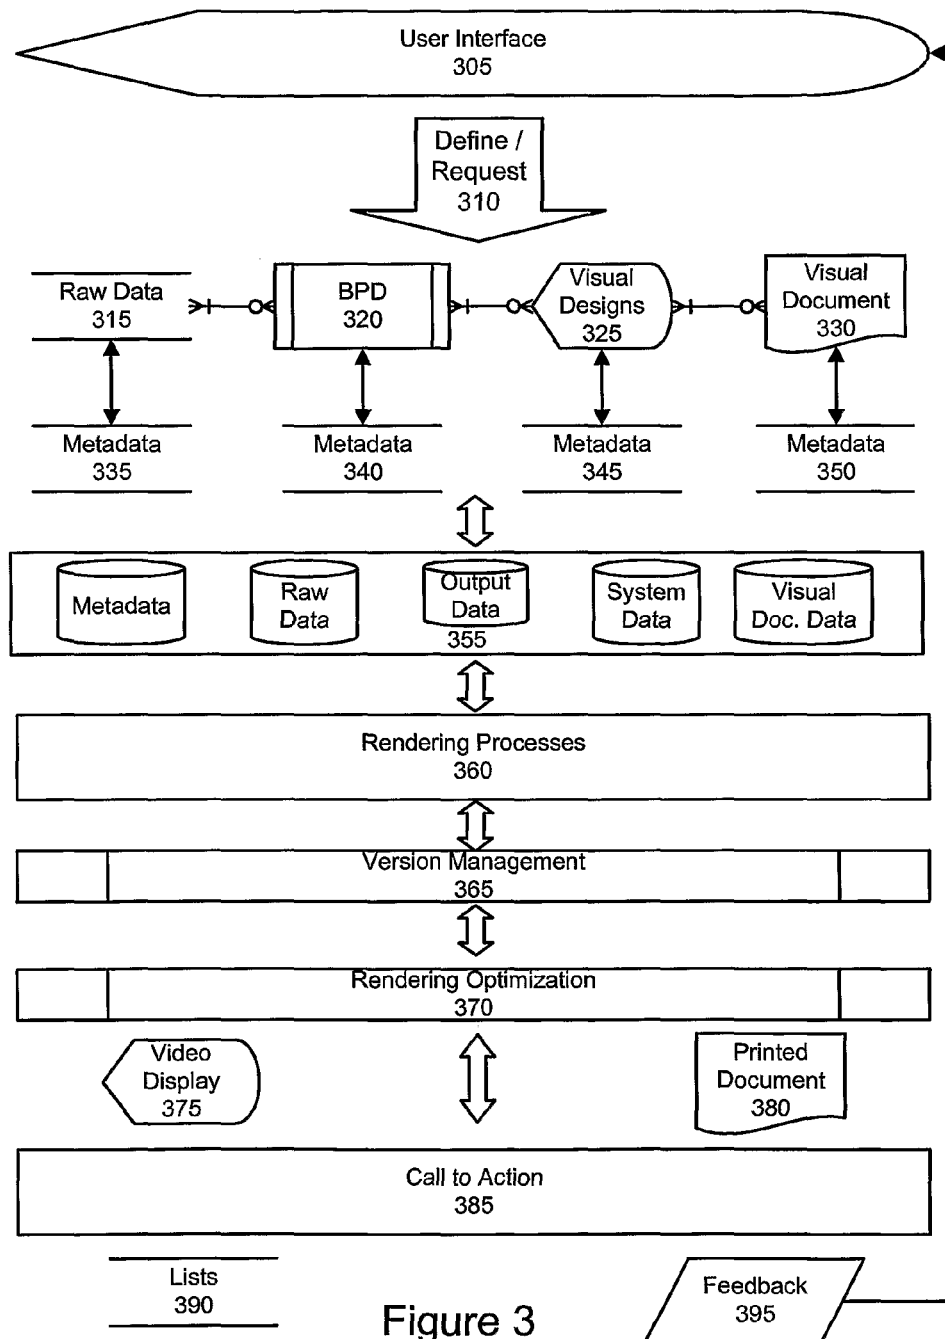


Figure 2B



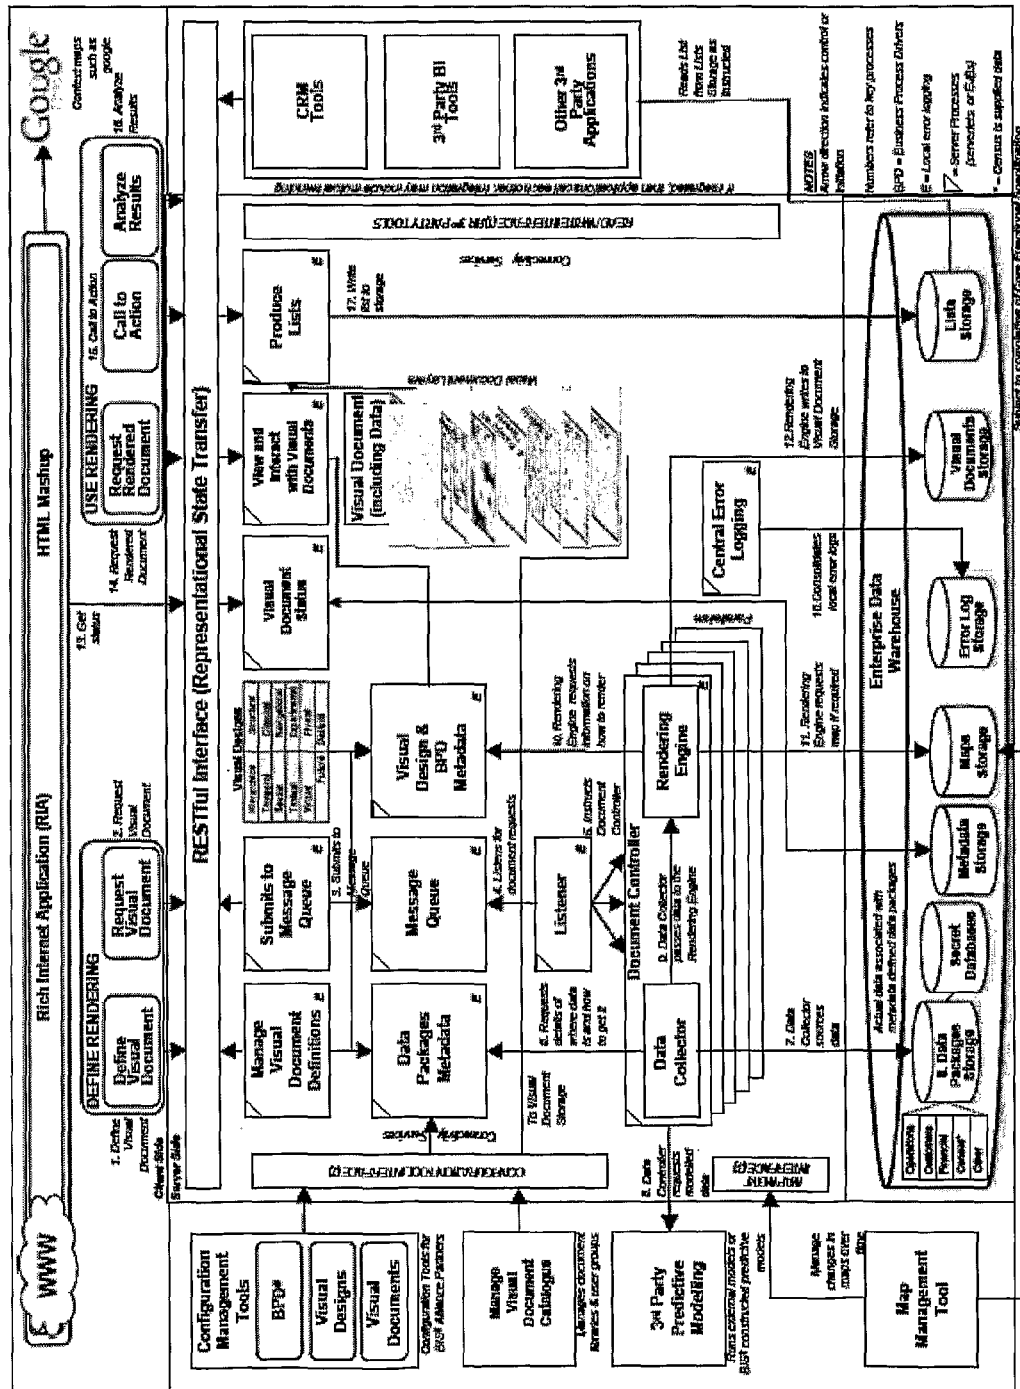


Figure 4




Figure 5

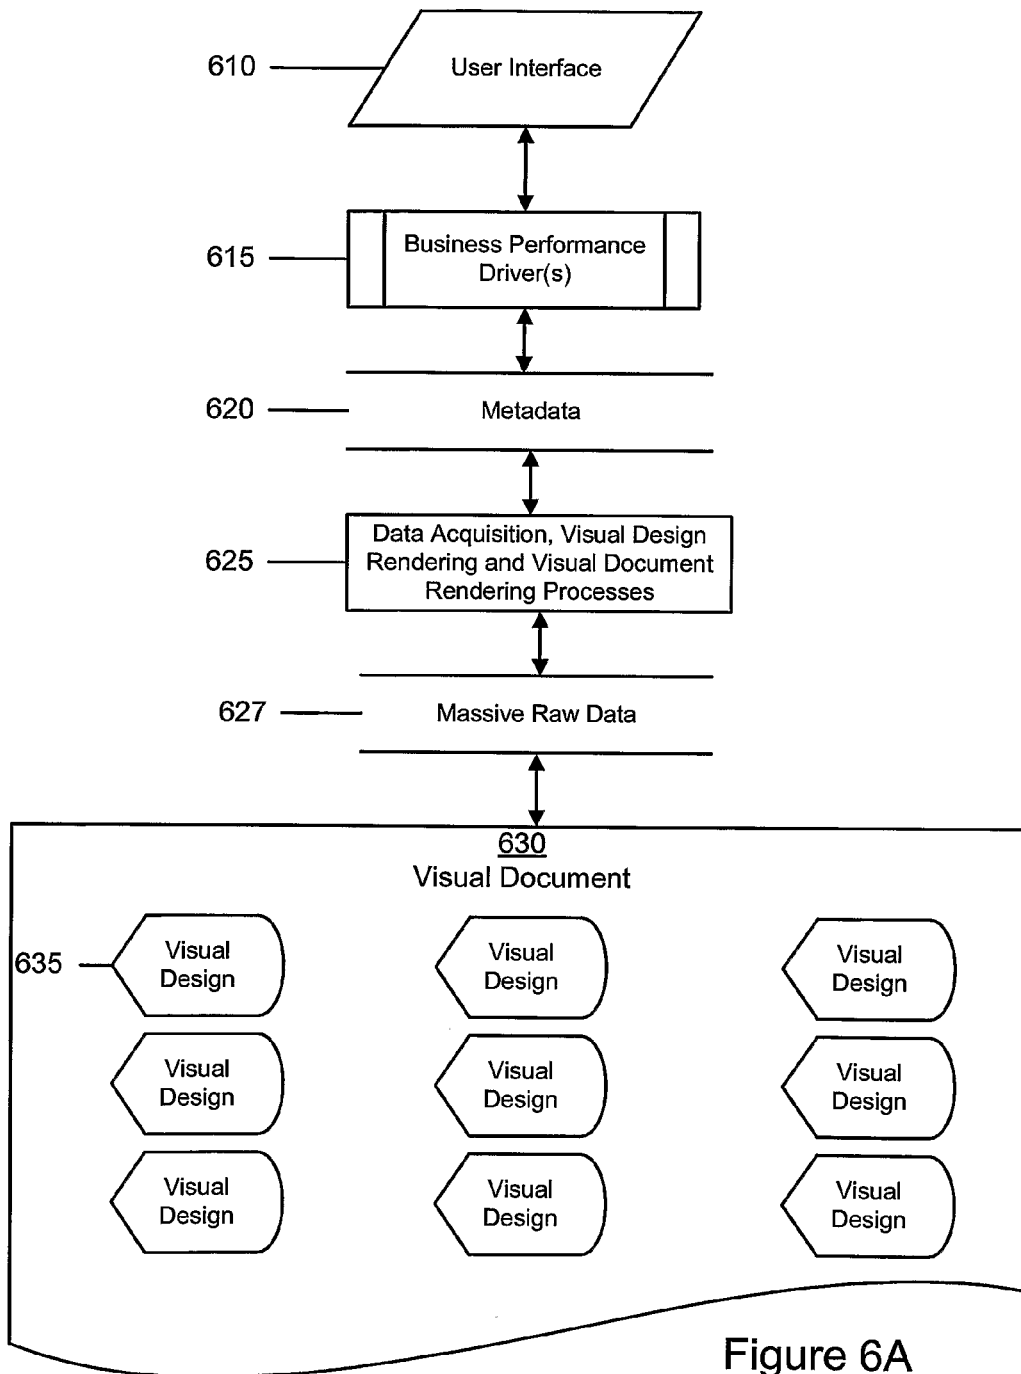


Figure 6A

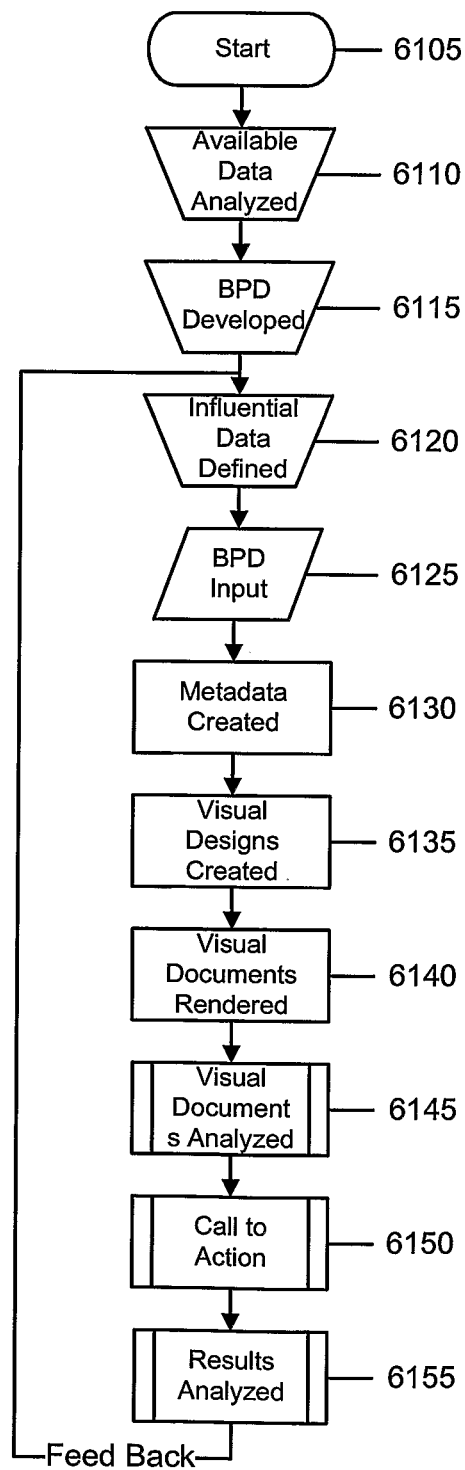


Figure 6B

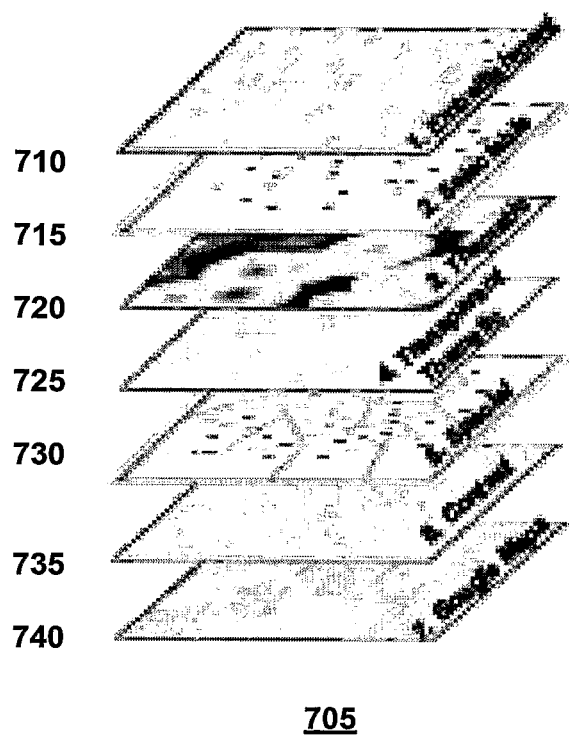


Figure 7

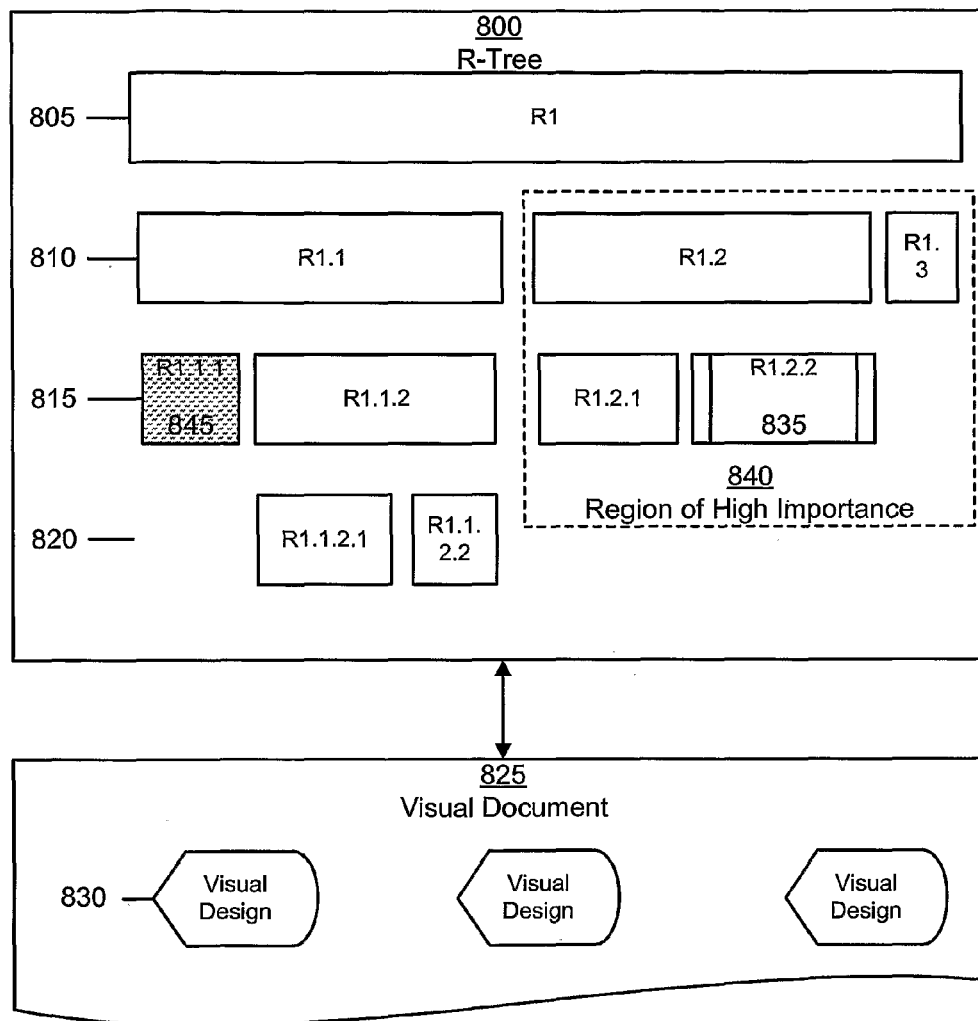


Figure 8




Figure 9

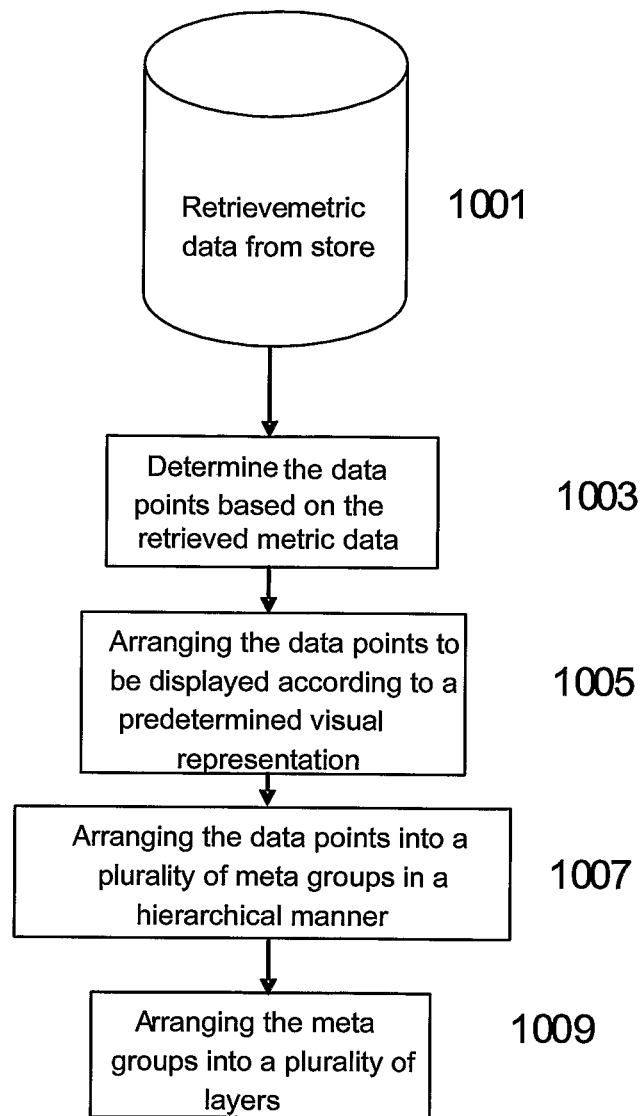


Figure 10




Figure 11

METHOD OF GRAPHICALLY REPRESENTING A TREE STRUCTURE

This application is a National Stage Application of PCT/NZ2009/000110, filed 18 Jun. 2009, which claims benefit of U.S. Ser. No. 61/074,347, filed 20 Jun. 2008 and U.S. Ser. No. 61/115,036, filed 15 Nov. 2008 and U.S. Ser. No. 61/140,556, filed 23 Dec. 2008 and which applications are incorporated herein by reference. To the extent appropriate, a claim of priority is made to each of the above disclosed applications.

FIELD OF THE INVENTION

The present invention relates to a method of graphically representing a tree structure.

BACKGROUND

A chart or graph is described in Wikipedia as a type of information graphic or graphic organizer that represents tabular numeric data and/or functions. Charts are often used to make it easier to understand large quantities of data and the relationship between different parts of the data. Charts can usually be read more quickly than the raw data that they come from. They are used in a wide variety of fields, and can be created by hand (often on graph paper) or by computer using a charting application.

Traditional charts use well established and often poorly implemented ways of representing data. Many tools exist to help the user construct very sophisticated representations of data but that sophistication typically results in less meaningful charts. Embodiments of the present invention aim to overcome this problem.

It is known to use charting wizards such as those that are available in Excel and various other systems such as those provided by, for example, IBM. In addition there are multiple Business Intelligence (BI) tools available to users to enable users to analyze data in an attempt to create meaningful feedback. However, as the amount of data increases, so does the complexity of the visual representations created by the analysis of the data. These complex representations can end up swamping parts of the visual representation that is most required and relevant to an end user.

Further, the focus of existing known methods of graphically representing data is on providing a single visual design, or type of visual or graphical representation, to represent data. That is, to produce, for example, a single bar graph to be displayed, or a single pie chart to be printed. This is very limiting to a user who may want to show various different aspects of the data in a single document.

Indexing methodologies, such as R-tree indexing for example, are used in conjunction with databases to categorize data and place the data in a hierarchical format. It is also known to use self organizing maps to visually represent data. However, self organizing maps can be very difficult and arduous to interpret. Also, it has not previously been known to use the indexing methodologies, in particular the R-tree indexing, as a display mechanism on its own. Further, no weight has been placed on how nodes in a tree are positioned relative to other nodes other than by the mere child parent relationship.

The present invention aims to overcome, or at least alleviate, some or all of the mentioned problems, or to at least provide the public with a useful choice.

SUMMARY OF THE INVENTION

Various concepts are herein disclosed as set out in the claims at the end of the specification.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

FIG. 1A shows a NASDAQ Heat Map Example;

FIG. 1B shows a NASDAQ Heat Map Intra Day Data Example;

FIG. 1C shows a diagrammatical representation of some key terms;

FIG. 2A shows a system concept diagram according to an embodiment of the present invention;

FIG. 2B shows an overview of the software modules in the described system.

FIG. 3 shows a general overview of the data flow within the system according to an embodiment of the present invention;

FIG. 4 shows an architectural overview of the described solution according to an embodiment of the present invention;

FIG. 5 shows a high-level system delivery overview of the described solution according to an embodiment of the present invention;

FIG. 6A shows a general data flow diagram according to an embodiment of the present invention;

FIG. 6B shows a flow diagram according to an embodiment of the present invention;

FIG. 7 shows the concept of layers according to an embodiment of the present invention;

FIG. 8 shows an R-tree representation according to an embodiment of the present invention;

FIG. 9 shows a conceptual block diagram of a system according to an embodiment of the present invention;

FIG. 10 shows a flow diagram of a method according to an embodiment of the present invention;

FIG. 11 shows how embodiments of the present invention may be incorporated within a gaming environment;

DETAILED DESCRIPTION OF THE INVENTION

The following described invention is suitable for use in conjunction with other methods, and the incorporation into one or more systems, for example as described in METHODS, APPARATUS AND SYSTEMS FOR DATA VISUALISATION AND RELATED APPLICATIONS (earlier filed by the applicant in the entirety as U.S. provisional patent application Ser. No. 61/074,347 filed on 20 Jun. 2008), which is incorporated by reference, and a portion of which herein follows.

Four key terms (or concepts) form the foundation of the specification set out in this document and accordingly have been defined as follows:

The four key terms are:

Business Performance Drivers (BPD)

BPD Packages

Visual Designs

Visual Documents

The key terms are defined as follows:

Business Performance Drivers (BPDs): A Business Performance Driver (BPD) is a business metric used to quantify a business objective. For example, turnover, sales. BPDs are Facts (sometimes referred to as measures). Facts are data items that can be counted. For example, Gross Sales; Units Sold. BPDs comprise of:

1. Measures: Data items that can be counted. For example, Gross Sales; Units Sold.

2. Dimensions: Data items that can be categorized. For example, Gender; Locations.

3

3. Restrictions can be applied to BPDs. These filter the data included. For example a restriction of "State="CA"" may be specified to only include data for California.

4. Normalizations can be applied to BPDs. These specify (or alter) the time period the BPD refers to. For example—Daily Units Sold, Monthly Profit. The combination of BPDs, Restrictions and Normalizations provides the flexibility to create many ways of looking at data without requiring extensive definition effort.

In other words a Business Performance Driver (BPD) is a 'measure' that can be normalized. Measures are data items that can be counted. For example, Gross Sales; Units Sold. BPDs might be displayed on visualizations. For example, Revenue earned per store on a map. Restrictions and/or Normalizations could be applied to a BPD. The following table provides examples of these:

Scenario	Business Example	
BPD (no normalization or restriction)	Revenue	20
BPD with restriction	Revenue earned in the state of California	
BPD with normalization	Revenue earned in week 1 of 2008	
BPD with restriction and normalization	Revenue earned in the state of California in week 1 of 2008	25

BPD Packages: A BPD Package is made up from a set of related BPDs. This relationship (between a BPD Package and its BPDs) is defined using metadata. BPD Packages can be thought of as the Visual Document's vocabulary.

Visual Designs: Visual Designs are a classification of the different types of visualizations that a user may choose. Within each Visual Design, there are a number of visualizations. For example, the 'spatial' category can have retail store location maps or geographical location maps. The software solution allows users to select one visualization (one visual form within a Visual Design category) to create a Visual Document.

Visual Document: A Visual Document contains visual representations of data. Access to the data used to construct the visual representation is in many ways analogous to a textual document. A Visual Document is constructed by applying BPD data to a specific Visual Design. It is designed to illustrate at least one specific point (using the visualization), supports the points made with empirical evidence, and may be extended to provide recommendations based on the points made. The Visual Document is a deliverable to the user.

Dimensions	Dimensions are data items that can be categorized. For example, Gender; Locations. Dimensions might be displayed on visualizations. For example product categories on a shop floor.	
Fact	See Business Performance Drivers (BPDs)	
Measure	See Business Performance Drivers (BPDs)	
Normalizations	Can be applied to BPDs. These specify (or alter) the time period the BPD refers to. For example - Daily Units Sold, Monthly Profit. The combination of BPDs, Restrictions and Normalizations provides the flexibility to create many ways of looking at data without requiring extensive definition effort. Refer to definition of BPDs for examples.	55
Restrictions	Can be applied to BPDs or Dimensions. These filter the data included. For example a restriction of "State = "CA"" may be specified to only include data for California. A BPD or Dimension could be restricted by Compound Statements (series of restrictions using AND/OR	60
		65

4

-continued

statements). For example, Revenue from all stores where state = California AND units sold >200 units. Restrictions have the following types:

Restriction Type	Definition	Example	Business Context
=	Equal to	State = 'CA'	Revenue earned within the state of California
>=	Greater than or equal to	Units Sold >=200	Revenue earned from stores where units sold were greater than (or equal to) 200 units
=<	Less than or equal to	Revenue <=\$50,000	Revenue earned from stores where Revenue was less than (or equal to) \$50,000
>	Greater than	Units Sold >200	Revenue earned from stores where the number of units sold were greater than 200 units
<	Less than	Units Sold <200	Revenue earned from stores where the number of units sold were less than 200 units
IN	In (list)	State IN ('CA', 'NY')	Revenue earned from stores within the states of California and New York
BETWEEN	Values between X and Y	Product Code between '124' and '256'	Revenue earned from product codes 124 to 256 (inclusive)
NOT =	Not Equal to	State NOT = CA	Revenue earned from stores outside the state of California.
NOT IN	Not in (list)	State NOT IN ('CA', 'NY')	Revenue earned from outside the states of California and New York.
NOT BETWEEN	Values not between X and Y	Store Code NOT Between 105 and 110	Revenue earned from stores excluding stores with a store code between 105 and 110 (inclusive).

Heatmaps: A heat map is a graphical representation of data where the values taken by a variable in a two-dimensional map are represented as colors. A very similar presentation form is a Tree map.

Heat maps are typically used in Molecular Biology to represent the level of expression of many genes across a number of comparable samples (e.g. cells in different states, samples from different patients) as they are obtained from DNA microarrays.

Heat maps are also used in places where the data is volatile and representation of this data as a heat map improves usability. For example, NASDAQ uses heat maps to show the NASDAQ-100 index volatility. Source: Wikipedia⁷

This is shown diagrammatically in FIG. 1A. Some blocks are colored green, which means the stock price is up and some blocks are colored red, which means the stock price is down. The blocks have a varying deepening of the relevant color to indicate the direction that the stock is moving. The deeper the color, the bigger the move.

If a user hovers over a stock, additional intra-day data is presented—as shown in FIG. 1B: Source: Nasdaq.com¹¹

The key terms are set out diagrammatically in FIG. 1C. Visual designs **110** are individual visualization techniques. One or more are applied to visualize BPD packages **115** to create visual documents **120**.

Many organizations are facing massive and increasing amounts of data to interpret, the need to make more complex decisions faster, and accordingly are turning to data visualization as a tool for transforming their data into a competitive advantage. This is particularly true for high-performance companies, but it also extends to any organization whose intellectual property exists in massive, growing data sets.

One objective of the described solution is to put experts' data visualization techniques in the customer's hands by skillfully guiding the end user through choosing the right parameters, to display the right data, and to create its most useful visualizations to improve business performance.

The described solution is a generic tool and can apply to multiple business areas that require decisions based on and understanding massive amounts of data. The resulting browser-based output is defined as a 'Visual Document'.

The solution provided is summarized in FIG. 2A.

The system identifies user tasks **201** in the form of defining visual documents, requesting visual documents, requesting rendered documents, calls to action, and analyzing results. These tasks are then detected by the system in conjunction with other systems **203**, which include CRM applications, third party Business Intelligence (BI) Tools and other third party applications, all of which may access data stored in an enterprise data warehouse (EDW). The visual design layer concept **207** may be utilized within the visual documents **205**. The creation of the visual documents is made in conjunction with a number of different defined visual design types **209**, BPD packages **211**, spatial analysis maps **213** and other application components **215**, such as application servers and application infrastructure.

A Visual Document contains visual representations of data. Access to the data used to construct the visual representation is in many ways analogous to a textual document. It is constructed by applying Business Performance Driver(s) (BPD) data to a specific Visual Design (Visual Designs are grouped into ten classifications).

A Visual Document is designed to illustrate at least one specific point (using the visualization), support the points made with empirical evidence, and may be extended to provide recommendations based on the points made. The Visual

Document is the actual deliverable from the software to the software user. Visual Documents may be stored, distributed or analyzed later, as needed.

The Visual Document is fed by data and a metadata database that stores definitions of BPDs—the BPDs are the focus of the Visual Document. A Business Performance Driver is a business metric used to quantify a business objective. Examples include, gross sales or units sold. For instance, the Visual Document may be used to graphically depict the relationship between several BPDs over time.

In the Visual Document, data is rendered in up to seven layers in one embodiment. However, it will be understood that the number of layers may be varied as needed by the user. Specific Visual Document Layers are described herein. However, it will be understood that further Visual Document Layers may be included over and above the specific types described.

Visual Designs are explicit techniques that facilitate analysis by quickly communicating sets of data (termed BPD Packages) related to BPDs. Once constructed, Visual Documents may be utilized to feed other systems within the enterprise (e.g., Customer Relationship Management (CRM) systems), or directly generate calls to action.

The described solution utilizes the best available technical underpinnings, tools, products and methods to actualize the availability of expert content.

At its foundation, the solution queries data from a high performance enterprise data warehouse characterized by parallel processing. This database can support both homogeneous (identical) and heterogeneous (differing but intersecting) databases. The system is adaptable for use with a plurality of third party database vendors.

A scalable advanced web server framework can be employed to provide the necessary services to run the application and deliver output over the web. A flexible and controllable graphics rendering engine can be used to maximize the quality and speed levels required to support both static and dynamic (which could be, for example, animated GIF, AVI or MPEG) displays. All components can operate with a robust operating system platform and within secure network architecture.

Pre-existing (and readily available) third party components can be employed to manage user security (e.g. operating system security), industry specific applications and OLAP (Online Analytical Processing) or other more traditional reporting. The described solution is designed to facilitate speedy and reliable interfaces to these products.

A predictive modeling interface assists the user in analyzing forecasted outcomes and in 'what if' analysis.

Strict security, testing, change and version control, and documentation standards can govern the development methodology.

Many organizations are facing massive and increasing amounts of data to interpret, the need to make more complex decisions faster, and accordingly are turning to data visualization as a tool for transforming their data into a competitive advantage. This is particularly true for high-performance companies, but it also extends to any organization whose intellectual property exists in massive, growing data sets.

This clash of (a) more data, (b) the increased complexity of decisions and (c) the need for faster decisions was recently recognized in an IDC White Paper (Gantz, John et. al.; IDC White Paper; "Taming Information Chaos: A State-of-the-Art Report on the Use of Business Intelligence for Decision Making" November 2007), which described this clash as the "Perfect Storm" and that this 'storm' will drive companies to make a quantum leap in their use of and sophistication in analytics.

Today's business tools and the way they operate barely allow business users to cope with historical internal data, let alone internal real time, predictive, and external data.

Hence, a new paradigm in business intelligence solutions is required.

System Overview

As explained above, FIG. 2A shows a high-level overview of the system.

There are five key components to the system. These are:

1. Visual Documents;
2. Visual Designs;
3. Business Performance Drivers (and BPD Packages);
4. Spatial Maps;
5. Application Components.

A description of each of these components is set out below under the respective headings.

Visual Documents

The Visual Documents form the core of the solution from a user perspective. This may include visualization(s), associated data and/or metadata (typically the visual form) that the user defines requests and interacts with. The Visual Documents may consist of single frames or animated frames (which could be, for example, implemented in AVI, GIF or MPEG format or a sequence of still images).

The Visual Document is typically viewed in a dynamic web browser view. In this interactive view the user may observe, select and navigate around the document.

Once created, the Visual Documents may be stored in the database and may be distributed to key persons (printed, emailed etc.) or stored for later use and analysis.

Visual Designs

The Visual Designs are a classification of the different types of visualizations that a user may choose. Within each Visual Design category, there are a number of visualizations. For example, the 'spatial' category can have retail store location maps, network maps or geographical location maps, such as, for example, maps available from Google™ or Yahoo™.

The described system allows users to select one or more visualizations (e.g. one visual form within a Visual Design category) to create a Visual Document.

There are ten Visual Design categories defined below, however it will be understood that further Visual Designs are envisaged, as well as the number of visualizations within each classification and the number of classifications.

Visual Designs are a classification of the different types of visualizations that a user may choose. Within each Visual Design, there are a number of visualizations.

For example, the 'spatial' category can have retail store location maps or geographical location maps.

The visual design types include:

- Hierarchical
- Temporal
- Spatial
- Textual
- Virtual
- Structural
- Classical
- Pivotal
- Navigational
- Interactive

1. Hierarchical Visual Designs

One purpose of a hierarchical visual design is to present large scale hierarchical data in one display. It is a picture for understanding, monitoring, exploring and analyzing hierarchical data.

Key elements of hierarchical visual designs are:

Data is hierarchical.

Structure of data can determine hierarchy.

They can be overlaid with connections.

This type of visualization may be automatically generated from a table of contents. This automatically generated hierarchy then becomes a special layer over which specific information can be overlaid.

The Hierarchical Visual Design is a hierarchical diagram such as an organizational chart or a correlation matrix.

This Visual Design has at least one natural centre and typically has a higher density toward the fringes of the visualization. The Hierarchical Visual Design can typically be considered as a 'tree' structure. The nodes and vertices within the tree structure are best if they are generated automatically from a dataset. This tree structure is a good example of a Special Layer.

The development process will include building a tree that is optimized for this type of Visual Design including heat mapping techniques.

Large scale hierarchical data is represented using various techniques such as mapping to icons, shapes, colors and heights.

Typical uses include mapping of web pages, organizational charts, decision trees and menu options.

2. Temporal Visual Designs

One purpose of a temporal visual design is to present temporal based data, such as, for example, revenue per day, in a specially designed calendar or time series view. This calendar view will enable users to view thematic layers that display BPD information such as revenue or sales.

This type of visual design is a completely data defined Visual Design. The key input values are typically 'start' and 'end' dates along with the 'number' of variables to be displayed.

The simplest, and potentially the most useful, Visual Design Special Layer may be a carefully drawn calendar. The calendar may then become a useful Visual Design for date-based Visual Documents.

Temporal analysis is one of the fundamental methods of almost all analysis. Using temporal high density visualizations, users will be able to overlay high density Thematic Layers on well designed Special Layers such as the spiral data visualization shown in the above examples. This analysis can be applied in everything from customer frequency and spend analysis to analysis of the impacts of time of day on the management of a mobile phone network.

It is considered that temporal design patterns are particularly important in terms of analytics as the majority of analytics are time based. Described herein are several examples of producing temporal visual designs.

Non Contiguous Time—For example, weekends can be represented in some interesting ways. The simplest way being not to show them.

Non-linear Time—This allows multiple years of history to be shown where the oldest data is spatially compressed in the Visual Design.

Temporal Special Layers—These can be used to compare quite disjointed types of data. For example, the relationship between external public events, operational payroll sizes and sales revenue. There exists no easy way to numerically join this data together, visually this data can be joined. The technique combines well with simple correlations as it is possible to combine these distinct datasets to show correlations.

Control—One important consideration in visualizing temporal data is the gaining of scientific control. For example, seasonal variables. This is particularly interesting as one year is always different from the next.

Quite simply, the start date of each year is never the same as the next, and moving external events such as Easter and ‘acts of God’ such as weather make precise comparison very difficult.

3. Spatial Visual Designs

One purpose of a spatial visual design is to present an overview of large scale numerical data in one spatial display (i.e. a space) for understanding, monitoring and analyzing the data in relation to a space.

This type of visual design combines together base maps provided by third parties with rendered thematic layers. These “mash-ups” are user definable and accessible to users.

For example, third party base maps may include customer-owned spatial maps or readily available base maps such as those provided by Google™ Maps or Yahoo™ Maps. The system provides powerful thematic layers over one of these spatial base maps.

One example of a spatial visual design is available at www.weather.comⁱⁱⁱ. This map shows two layers—(1) an underlying heat map overlaid with (2) actual temperature at specific cities. The points are useful as the state boundaries allow the user to determine with relative ease which city is being referenced. The underlying heat map is useful as it allows the user to see the overall trend at a glance.

A second example is available at Information Aesthetics^{iv}. This example shows the travel time from the centre of London outwards using various methods of travel. The use of heat maps here shows very clearly the relationship between distance from the centre of London and travel time.

In a further example, the ‘spatial’ category of visual design can have retail store location maps, network maps or geographical location maps, such as, for example, maps available from Google™ or Yahoo™.

Numerical data may be independently mapped using parameters such as hue, saturation, brightness, opacity and size distributed across a defined geographical space.

Geographic mapping has a wide range of uses. In fact with the wide availability of high quality base maps, the world is becoming spatially enabled. Mapping applications can be used for a huge variety of tasks, from customer relationship management to drive time analysis, site selection to insurance risk analysis and telecommunications network analysis.

4. Textual Visual Designs

One purpose of textual visual designs is to enable business users to interact and query seamlessly from the structured to the unstructured world.

While it is possible to do basic numeric analysis on variables such as hit frequency and number of clicks per hour, the key method is to use a special layer to construct a sensible schematic of the unstructured data then overlay BPDs. Simply put, the described solution will leverage information visualization to bring structure to the unstructured world.

For example, a heat map may be used as part of a textual visual design.

Unstructured textual information is a huge area of growth in data storage and intuitively, the business intelligence industry expects this data to become a valuable asset. The described solution provides information visualization capabilities that overlay and draw out the non-numeric, but actionable, observations relating to unstructured data, in order to link the numeric data warehouse to the unstructured world.

There are a multitude of Special Layers that may be used with textual data. These textual Special Layers extend from building self organizing maps of textual information to diagrams showing the syntax hierarchy of the words used in a document.

A self organizing map (SOM) consists of components called nodes or neurons. Associated with each node is a weight vector of the same dimension as the input data vectors and a position in the map space. The usual arrangement of nodes is a regular spacing in a hexagonal or rectangular grid. The self-organizing map describes a mapping from a higher dimensional input space to a lower dimensional map space. The procedure for placing a vector from data space onto the map is to find the node with the closest weight vector to the vector taken from data space and to assign the map coordinates of this node to our vector—Source: WikipediaError! Bookmark not defined.

5. Virtual Visual Designs

One example of a virtual visual design is a 3D representation of a virtual environment. 3D worlds generate far more accurate and complete data than the real world. As these 3D worlds grow in popularity and become more immersive, the potential for business intelligence tools to be applied to this environment grows significantly.

One example application of the use of a virtual visual design is a retail space analysis tool where transaction data is under-laid as the color of the carpet or shelves. In the case of the shelves, the shelves can also show representations of the products on the shelves.

6. Structural Visual Designs

One purpose of a structural visualization is to illustrate the structure of the data. For example, network topology or interconnection between data elements. The interconnections in the examples below show how a simple Special Layer construct can be used to illustrate quite complex connections.

One example of a structural type visual representation is that of the London underground map. The London underground map is a key historic map showing the schematic topology of the London underground. Using this map travelers can intuitively plan out complex routes and interconnects. Without this visualization, navigating the London underground system would be significantly more difficult and complex to understand.

These structural visualizations are very powerful and are closely related to spatial visualizations. Most of the thematic treatments that can be applied to a spatial visualization are equally applicable to a structural visualization.

Examples of uses for such a visual design type would be for visualizing call routing across a network, electricity grid management and route optimization.

It will be understood that a wide variety of Special Layers may be created in this space. These Special Layers essentially generate the structural schematic from the base data.

Typically the interconnections between nodes are used to generate the structure. One important aspect of the structural Special Layer is building the structure in such a way that interconnect line crossing is minimized.

7. Classical Visual Designs

Traditional charts provide a simple, common and well-established way of presenting data using classical visual designs. However, traditional charts are user-skill dependent and the herein described system may be used to apply guided Visual Design techniques to traditional charts to significantly extend their usefulness.

One example would be to show a line chart of Speed Vs Time in a simple two dimensional line graph. This type of basic graph shows the data clearly and allows the user to observe any geometric trends.

Some common charts that fall into this design category are as follows:

Scatterplots—Are Cartesian coordinates to show the relation of two or more quantitative variables.

11

Histograms—Typically show the quantity of points that fall within various numeric ranges (or bins).

Bar graphs—Use bars to show frequencies or values for different categories.

Pie charts—Show percentage values as a slice of a pie.

Line charts—Are a two-dimensional scatterplot of ordered observations where the observations are connected following their order.

8. Pivotal or Quartal Visual Designs

Different visualization methods have been suggested for high-dimensional data. Most of these methods use latent variables (such as principal components) to reduce the dimensionality of the data to 2 or 3 before plotting the data. One problem with this approach is that the latent variables sometimes are hard to understand in terms of the original variables.

The parallel coordinate (PC) scheme due to Inselberg and others attempts to plot multivariate data in a completely different manner. Since plotting more than 3 orthogonal axis is impossible, parallel coordinate schemes plot all the axes parallel to each other in a plane. Squashing the space in this manner does not destroy too much of the geometric structure. The geometric structure is however projected in such a fashion that most geometric intuition has to be relearned, this is a significant drawback, particularly for visualization of business data.

Pivotal or Quartal visual designs allow the user to display higher dimensional data in a lower dimensional plot by ranking and splitting variables across various axes. This method may for example be used to display 3D data in a 2D plot.

9. Navigation Visual Design

Navigational visualizations use a highly visual interface to navigate through data while maintaining the general context of the data. This data visualization method may use other visual design types so it is differentiated more by the style of how it is used than the implementation standard.

Photosynth for example is a powerful navigational tool for moving between images, its display is designed for navigation of large numbers of linked images.

One illustrative navigational representation example is shown by Ubrowser. This navigational visualization example shows web pages represented in a geometry design. The web pages can be navigated through by spinning the cube shown in the example.

Navigational visualizations are designed for users to interactively move through the data. The objective of the visualization is to present a large volume of data in such a way as to enable users to move through the information and gain an understanding of how the data links together.

A number of display techniques are known for displaying information with regard to a reference image (the combination referred to as primary information). Where the limit of primary information is reached a user may wish to know more but be unable to further explore relevant information. A user may also simply wish to explore other aspects although there is more primary information to explore.

A key element of navigational visual designs is that they are interactive and are designed to assist in data navigation and data way-finding rather than for analytical purposes.

10. Interactive Visual Designs

This classification is for significantly advanced or interactive visual designs which do not fit within the preceding classifications.

These visualizations vary in nature from pure abstract forms to more tangible forms of visualizations. The key difference is that these visualizations may not be classified within the preceding Visual Design classifications due to their advanced nature or interactivity.

12

Any Visual Design layer considerations will be dependent on the interaction being considered.

There is opportunity to use common associations to provide iconic views of key events; the common associations are created using the interactive tools and asking users for feedback on the relevant icons. This feedback is then developed into a learned interactive system to provide iconic data representations.

Eye movement sensors can be used to control the interactivity and to learn information about relevant icon usage and control interactivity.

A wide range of user interfaces are used in conjunction with computer systems. Generally these are simply used to provide command or data inputs rather than to analyze the underlying behavior of a user in the context of the operation of a software application.

It would be desirable to operate software applications running on a computer on the basis of observed user behavior in the context of a software application.

Business Performance Drivers (and BPD Packages)

Business Performance Drivers (BPDs) are a metric applied to data to indicate a meaningful measurement within a business area, process or result. BPDs may be absolute or relative in their form of measurement.

The Business Performance Driver (BPD) concept differs from the known KPI concept by introducing BPDs that

- (1) may have multiple dimensions,
- (2) place the BPD in the context of the factors used to calculate them,
- (3) provide well understood points of reference or metadata around which visual document creation decisions can be made, and
- (4) may contain one or more methods of normalization of data.

Common groups of BPDs are called BPD Packages. For example, BPDs relating to one industry (say, telecommunications) can be grouped into one BPD Package. BPDs may be classified into one or more BPD Packages. For example, Net Revenue with normalizations available of per customer or per month may be applicable in a number of industries and hence, applicable to a number of BPD Packages.

Spatial Maps

Spatial maps allow for a user-owned and defined spatial map and/or for the user to use publicly available context maps such as Google™ Maps or YahooSM Maps. In either case, the user can display selected BPDs on the chosen spatial map.

Typically, a user-owned spatial map may be the inside floor space of a business and a publically available context map may be used for displaying BPDs on a geographic region e.g. a city, county, state, country or the world.

Application Components

The described application includes two main components, the Application Servers and the Application Infrastructure.

The Application Server includes a number of servers (or server processes) that include the Rendering Engine (to make (or render) the Visual Documents), Metadata Servers (for the BPD Packages, the Visual Designs and the BPDs) and the Request Queue.

The Application Infrastructure is also comprised of a number of servers (or server processes) that may include a Listener (which 'listens' for document requests) and central error logging.

Based on the user selections made above (Visual Documents, Visual Designs and BPDs), the user can click on an action and send a communication to a third party system (CRM, Business Intelligence or other application). The third

13

party system could, for example, load the list from the solution and then send out a personalized email to all members on that list.

According to one embodiment, the described server components of the application are a Java based application and utilize application framework such as the IBM™ WebSphere application server framework, other platforms and server applications may be utilized as alternatives. The client application may be a mashup that utilizes the server components or it could be a rich internet application written using the Adobe™ Flash framework.

Other key elements of the system may include:

Parallelism—Parallel processing to increase responsiveness or to increase workload scalability of queries or Visual Documents. This parallelism may also decrease response time for larger visual documents in particular animated images may be executed in a parallel fashion.

Security—System and user-access security. This security may be a combination of authorization and authentication. The security framework may be implemented using the application framework.

Map Updates—A map management tool to update user-owned spatial maps.

14

Predictive Modeling—This may be an interface to third-party predictive models.

Configuration Tools—The application may be supported by configuration tools to enable rapid deployment of the application.

Modular Overview

Module Descriptions

The diagram shown in FIG. 2B shows an overview of the software modules in the described system.

These modules are described in the subsequent table. More detailed descriptions and diagrams of each of the software modules are provided below.

The table below outlines the following four items in relation to each module:

1. Technology System Component: This is the name given to the system component; this name matches the name in the above diagram.
2. High Level Functional Description: Describes the role of the software module.
3. Caching: Indicates whether this module uses caching to optimize performance.

Technology System Component	High Level Functional Description	Caching
1. Rendering Engine	Produces images and animations; could use Google™ Maps or Yahoo™ Maps for spatial context map. The development of Special Layers enables Visual Document produced to have unique capabilities that were not previously readily available.	Yes
2. Parallelism Engine	Enables parallel execution of requests for high volume of Visual Document output and rapid results delivery to users. The preferred application framework selected is the IBM™ WebSphere product. This framework enables the application to be scaled across multiple servers.	Yes
3. Map Management Tool	Provides key map editing features (specifically CAD like) and map version control (desktop and enterprise) tools.	Yes
4. OLAP Reporting	Industry standard online analytical reporting. For example, sorting, filtering, charting and multi-dimensional analysis. It is desirable that the user interaction with the data selection process in the data view is seamless with the data visualization view. For example, if the user selects 5 customers from the data view, the same 5 customers should be selected in the visualization view. This means that the solution may be a hybrid view (as discussed later). This hybrid view is a 'simple' view and is an interface to an industry leading OLAP tool. One option includes interfacing to the OLAP tool via a JDBC interface from the described solution or a web service model for the selection management.	Yes
5. Predictive Modeling System	An interface to external predictive modeling engines; may also have some modeling systems. For example, Self Organizing Maps (SOM).	Yes
6. Visual Design Management System	Tools for users to manage the different Visual Designs.	No
7. BPD Management and Data Access System	Tools for users to manage the different BPD Packages and their associated BPDs. Contains Data Access capability that enables data to be queried from RDBMS (or potentially other data sources).	No
8. Output Management System	For management of the documents (Visual Documents) within the system.	Yes

Technology System Component	High Level Functional Description	Caching
9. Infrastructure	Core system management functions including system logging and Request Queue management. The Request Queue is also described under parallelism and there may be crossover between these two module descriptions.	Yes
10. Security	Enables access to the system (or parts thereof) to be properly controlled and administered.	No
11. Interfaces	Allows services to be called by (or to call) external applications.	No
12. Implementation Tools	Tools to deploy and configure the software system.	Yes

Architectural Views of the System

This section contains descriptions and diagrams of the architectural views of the system. The architecture shows how the system components fit and operate together to create an operational system. If compared to a vehicle, the wiring diagrams, the physical body, the driving circle and key complex components like the engine would be shown in architectural views.

This view does not describe how the system is written; it describes the high-level architectural considerations.

Architectural considerations are typically implemented by one or more software modules. The modular view described herein lays out a high-level view of how the software modules are arranged.

FIG. 3 shows a general overview of the data flow within the system.

FIG. 4 shows the architectural overview of the described solution. This diagram is elaborated by the diagrams and descriptions in following sections of this document.

The following modules or components are shown:

Web interface Module **4105**: User interfaces are browser based or may be a web services client, a rich internet application or may be a thick client. In all cases the user interface uses the same interface to the back end services.

Rendering Definition Module **4110**: The user interface is used to define and request the rendering of Visual Documents

Rendering Use Module **4115**: Visual Documents are used for analysis, and precipitate calls to action.

Connectivity Services Module **4120**: The definition and rendering of Visual Documents is performed through a set of programs or services called the Connectivity Services.

Configuration Management Tools Module **4125**: Multiple versions of the basic elements; BPD, Visual Design, Visual Documents; are managed by a set of programs called the Configuration Management Tools.

Visual Document Management Catalog **4130**: One such Configuration Management Tool (**4125**) is a set of programs that manage a users' catalog of available Visual Documents.

Predictive Modeling Module **4135**: Predictive modeling is used for forecasting unknown data elements. These forecasts are used to predict future events and provide estimates for missing data.

Map Management Tool **4140**: Another of the Configuration Management Tools (**21125**) is the Map Management Tool. It is designed to manage versions of the spatial elements of a visual design such as a geographic map or floor plan.

Visual Document Definitions Management Module **4145**: Visual Document Definitions are managed through the use of metadata (**4175**).

Message Queue Submission Module **4150**: Requests for Visual Documents are handled through queued messages sent between and within processes.

Visual Design Type Module **4155**: Visual Documents are comprised of one or many Visual Designs in these categories.

Visual Document Status Module **4160**: The status of Visual Documents is discerned from the metadata and displayed on the user interface.

Interaction and Visual Document View Module **4165**: The user interacts with the Visual Documents through the user interface, and appropriate changes to and requests to read are made to the metadata.

List Production Module **4170**: Where additional output such as customer lists are required, they are requested using the user interface and stored in the EDW (**4215**).

Data Packages Metadata Module **4175**: Metadata is used to describe and process raw data (data packages).

Message Queue Module **4180**: Messages may be queued while awaiting processing (**4150**).

Visual Design and BPD Metadata Module **4185**: Metadata is used to describe and process the BPD's and Visual Designs associated with a particular Visual Document.

Visual Documents Module **4190**: Visual Documents may be comprised of layered Visual Designs.

Third Party Modules **4195**: Visual Documents may be used with or interact with other third party tools.

Listener Module **4200**: The listener processes messages (**4150**) in the message queue (**4180**)

Document Controller Module **4205**: The document controller is used to provide processed data to the rendering or query engines.

Central Error Logging Module **4210**: System errors are detected and logged in the EWP (**4215**).

EDW **4215**: All data is typically stored on a database, typically, multiple fault tolerant processors in an Enterprise Data Warehouse.

The following architectural components are described in more detail.

Architectural Component	Description
Connectivity Services	This is a common communication service that is used when sending messages between systems (i.e. the described solution and 3 rd party tools) and between the described application layer and the user interface layer.
Configuration Management Tools	Allows specialized users to configure Visual Designs and Visual Documents to their needs - which differ from the default configuration provided.
Manage Visual Document Catalog	Gives selected users the ability to search, sort, group, and delete Visual Documents in the Visual Document Catalog.
Predictive Modeling	External modeling systems that use data sent from the described solution to perform complex calculations to produce predictive data. This predicted data is piped through the described solution to the user.
Map Management Tool	This is an application that enables users to create modify and delete individual maps to manage the complete sequences, this is very appropriate for management of floor plans.
Data Packages Metadata	The services responsible for providing metadata that enables the requester (typically, Data Collector) to source the data for the BPD.
Visual Design & BPD Metadata	The services responsible for providing the metadata to the requester (typically the Rendering Engine) that enables the construction of the Visual Documents.
Request Queue	The Request Queue manages the communication of requests for rendering of Visual Documents. These communications may be scheduled.
Document Controller	The Document Controller consists of two components. The first is the Data Collector responsible for reading the appropriate metadata and retrieving the data from the EDW (Enterprise Data Warehouse). This data is passed to the Rendering Engine that is responsible for producing the Visual Document. Document Controllers run parallel Visual Document requests, build and store documents.
Read/Write Interface for 3 rd Party Tools	The described solution provides a common interface for 3 rd party tools to communicate with e.g. CRM applications.
3 rd Party BI Tools	One of the 3 rd party tools that the described solution may integrate with is an external OLAP tool.
Secret Databases	Secret databases are a method of sharing encrypted databases and providing a SQL interface that enables end users to run queries against atomic data without discovering the details of the data.

The following terms have been also been used in FIG. 4.⁴⁰ These are explained in more detail below.

-continued

Architectural Component	Description
Logging	Logging (for example, error logging and access logging) is an inherently difficult activity in a parallel independent and predominantly stateless system. The main issue that arises is that logging presents potential links between systems and therefore dependencies. Typically within the application, each server will be responsible for its own logging. This ensures that the system scales without degradation in performance. A separate process (central log reader) may be used to consolidate these logs dynamically as and when required.
Web Server	Web Servers respond to requests from users to provide Visual Documents. They read any required information from the metadata servers and Visual Document storage servers. If necessary they write Visual Document requests to the Request Queue.
Metadata Servers/Storage	Metadata servers are responsible for storage and user views of metadata. The metadata servers are also responsible for the validation of user rights to read Visual Documents (within the application).
Visual Document Storage	The Visual Document Catalog is a secure storage for all Visual Documents. Access is only possible when security requirements are met.
45 Data Collector	Typically the data collector queries the customer's data warehouse. The data warehouse can be augmented with additional subscribed embellishment data. This will provide the raw data that is represented visually back to the user.
50 BPD Packages Metadata	The described solution will use metadata to define groups of BPDs. These groups of BPDs are called BPD Packages. BPD Packages enable both internal data measures to be efficiently installed and external datasets to be provided. BPD packages contain no data.
55	A further high-level system delivery overview of the solution is set out as shown in FIG. 5.
60	The described solution 500 is hosted by the enterprise 510. The figure shows the logical flow from the submission of a request to the end result, viewing the rendered Visual Document.
65	The data being visualized belongs to the customer 512 and the submitted request is unknown to the entity running the visualization system 500.
	The controlling entity, integrators and customers may wish to have summaries of technical performance data (usage patterns, errors etc) sent from the operational system back to the integrator or controlling entity.

The system **500** has access to the data in a EDW **505**. The system utilizes a request queue **515** to control requests from a corporate network **510**. These requests are forwarded to a document controller **520**. The document controller **520** accesses both the EDW **505** and reads visual designs and BPD metadata services **525**, as well as data packages metadata services **530**.

The system described thus enables various methods to be performed. For example, data is transformed into visually interpretable information. The visually interpretable information is in the form of visual representations that are placed within one or more visual documents.

FIG. 6A shows a general data flow diagram for the described system.

The User Interface **610** allows the user to define BPD's **615** in terms of raw data **627**, which become the focus of the Visual Document **630**.

Further, the User Interface **610** allows the user, through automated expert help, to create the Metadata **620**, the most appropriate Visual Designs **635** that make up the Visual Document **625** in order to provide detailed analysis of data related to the BPD **615**. The data acquisition, visual design rendering and visual document rendering processes utilize massive amounts of raw data **627**.

The Metadata **620** is used by the Processes **625** to optimize the acquisition of the appropriate Data **627**, processing of the data into useful information, and to optimize the creation and rendering of the Visual Designs **635** and the Visual Document **630** that contains them.

This method includes the steps of providing comprehensive yet easy to understand instructions to an end user that has accessed the system and the visual design application. The instructions assist the end user in obtaining data associated with a theme, wherein the theme may be focused on objectives that have been derived from the data. The objectives may be business objectives, for example. In this way, the system guides a user carefully through the many choices that are available to them in creating the visual representations, and the system automatically tailors its instructions according to not only what the user requires, but also according to the data that is to be represented. The system focuses on providing instructions to enable a visual representation to be created that will enable an end user to more effectively understand the data that has been collated.

Further, the instructions assist the end user in determining one or more summaries of the obtained data that enable the end user to understand the theme, as well as organizing the determined summaries into one or more contextual representations that contribute to the end user's understanding of the theme.

Further, instructions are provided that assist an end user in constructing one or more graphical representations of the data, where each graphical representation is of a predefined type, as discussed in more detail below, and includes multiple layers of elements that contribute to the end user's understanding of the theme.

Finally, instructions are provided to assist an end user in arranging the produced multiple graphical representations in a manner that enables the end user to understand and focus on the theme being represented as well as to display or print the organized graphical representations. The system assists in the organization or arrangement of the representations, elements thereof, within the visual document so as to ensure certain criteria are met, such as, for example, providing a suitable representation in the space available, using the minimum amount or volume of ink to create the representation, and

providing a suitable representation that depicts the theme in a succinct manner, or visually simplistic manner.

The data being processed to create the graphical representations may be particularly relevant to the theme being displayed, disparate information or indeed a combination of relevant and disparate information.

There are multiple types of graphical representations that may be included within the visual document. The types are discussed in more detail below and include a hierarchical type, a spatial type, a virtual type, a classical type, a navigational type, a temporal type, a textual type, a structural type, a pivotal type, and an interactive type.

Further, the instructions may assist an end user in arranging the graphical representations in order to display high density data in a manner that conveys important information about the data, rather than swamping the end user with multiple representations that look impressive but do not convey much information.

In addition instructions may be provided to assist the end user in arranging the graphical representations to allow supplementary information to be added, where the supplementary information may be provided in any suitable form. Particular examples provided below depict the supplementary information being provided in subsequent visual layers that overlay the graphical representation. Alternatively, or in addition, supplementary information may include additional elements to be displayed within a single layer of the representation, for example, in the form of widgets.

FIG. 6B shows a flow diagram according to this embodiment of the invention.

Step **6105**: Process Starts. User decides to manage the business.

Step **6110**: Available data is identified and analyzed.

Step **6115**: Business Process Drivers (metrics defined in terms of the data to indicate a meaningful measurement within a business area, process or result).

Step **6120**: Data influencing the BPD metrics are identified.

Step **6125**: BPD's are input into a computer system

Step **6130**: BPD is categorized and appropriate metadata describing it is generated.

Step **6135**: Visual Designs to display the influential data are created.

Step **6140**: Visual Designs are aggregated into Visual Documents and rendered. Adjustments are made based on the freshness of all components (e.g., BPD, available data).

Step **6145**: Visual documents are analyzed by the end user.

Step **6150**: The end user decides on and implements actions based on the analysis in **6145**.

As touched on above, business performance drivers (BPDs) are used to enable more efficient data analysis so as to produce accurate and relevant visual representations of the data. A BPD is a form of advanced business measure wherein additional information is included, within the BPD that enables the system using the BPD to understand how to manipulate the BPD. That is, one or more intelligent attributes are included with the business measure to form the BPD, where those attributes reference or include information on how the BPD is to be processed or displayed. The form of processing and display may also be varied according to the device type or media upon which the business measures are to be displayed.

The attributes are attached to the business measure by storing the BPD in the form of a mark up language, such as, for example, HTML or XML. It will however be understood that any other suitable format for storing the BPD may be used where the attributes can be linked to the business measure.

In the example of HTML, the attribute is included as a tag. One such example would be to include the data or business measure within the body of the HTML code and follow the business measure with a tag that references the attributes, or dimensions, associated with that business measure.

Further, the attributes may also be modified or deleted, or indeed new attributes added, during or after the processing of the BPD so that the attributes are maintained, or kept up to date, bearing in mind the requirements of the entity using the BPD to visualize their data.

The business performance drivers, or measurable business objectives, are identified in order to create graphical representations of the business objectives, where those representations are placed within a visual document. A business objective may be, for example, a metric associated with a business.

Instructions are provided by the system to the end user, in order to assist the end user in establishing multiple business objectives as functions of available metrics, as well as assisting the user in organizing the business objectives into a contextual form that contributes to the end user's understanding of the business objectives.

Further, instructions are provided to assist the end user in constructing one or more graphical representations of the business objectives, where each graphical representation is of a predefined type, as mentioned above and described in more detail below. Further, each graphical representation includes multiple layers of elements that contribute to the end user's understanding of the business objective.

The elements within the graphical representation may include, for example, a shape, position, color, size, or animation of a particular object.

Instructions are also provided by the system to assist the user in arranging multiple graphical representations in a suitable manner that enables the end user to understand and focus on the business objectives being represented.

Finally, the end user is also assisted with instructions on how to display the organized graphical representations.

The following section describes a method of creating a visual representation of data in the form of a visual design.

The method includes the steps of the system providing instructions to an end user to assist the end user in constructing multiple graphical representations of data, where each graphical representation is one of a predefined type, as defined above and explained in more detail below, and the graphical representation includes multiple layers of elements that contribute to the end user's understanding of the data.

The system also provides instructions to an end user that assist the end user with arranging multiple graphical representations of different types within the visual representation in a manner that enables the end user to understand and focus on the data being represented, as well as providing instructions to assist the end user in displaying the visual representation in a suitable manner.

The visual representation may be displayed in a number of different ways, such as on a color video screen or a printed page. The information that is forwarded to the display device to create the visual representation may differ according to the type of display device so that the visual representation is produced in the best known suitable manner utilizing the advantages of the display device, and avoiding any disadvantages.

The data being displayed may be based on a measured metric or an underlying factor that affects a metric.

The elements within the graphical representation may include a shape, position, color, size or animation of a particular object.

Although a single visual document may include only one type of graphical representation, either in the form of multiple graphical representations or a single representation, there will also be situations where multiple types of graphical representations may be organized within a single visual document in order to convey different aspects of the data, such as, for example, temporal as well as spatial information. The inclusion of different types of graphical representations within a single document can provide an end user with a better understanding of the data being visualized.

Further, the single visual representation may be arranged to be displayed as an image on a single page or screen. This may be particularly useful where space is at a premium yet the user requires the visual representation to be provided in a succinct manner. For example, the user may request certain information to be displayed in a visual representation on a single mobile telephone display, or a single screen of a computer display, in order to show a customer or colleague the results of a particular analysis without the need to flick between multiple screens which can result in confusion, a waste of energy and ultimately a loss of understanding of the visual representations.

The same issue applies to printed representations, where the result of the system enabling a user to arrange a single representation, which may include multiple elements or layers, on a single page not only succinctly represents the data being analyzed but also saves the amount of paper being printed on and the amount of ink being used to print the document.

Further, the amount of ink required for a visual representation may be further reduced by providing instructions to the end user in a manner that directs them to control and use white space in a representation in an efficient manner so as to reduce the requirement of ink.

Multiple types of graphical representations may be merged together within a single visual document, or representation.

As mentioned above, instructions can be provided by the system to assist the end user in adding supplementary information to the visual representation, and the supplementary information may be provided in layers within the representation.

Visualization Framework

The following description provides the visualization framework that will support embodiments of the present invention. The description includes an overview of the importance of Visual Design including a brief historical recount of a world-recognized leading visualization. The description also sets out the Visual Design classifications for the described solution.

It will be understood that the Visual Design examples described in this section, are examples for illustrative purposes to identify the concepts behind how the visualization is produced. Therefore, it will further be understood that the concepts described can produce visual designs different to those specifically described. The Visual Design examples shown are also used to help the reader understand the narrative describing the Visual Designs.

The system described is specifically adapted to create actual specific visualization designs relevant to selected vertical and horizontal industry applications being deployed.

A vertical industry application is one that is associated with a solution directed at a specific industry, such as, for example, the entertainment industry. In this example, BPDs relevant to that industry are created, such as rental patterns of movies over different seasons.

A horizontal industry application is one that is associated with solutions across multiple industries. For example, the

BPD may be based on CRM analytics, which applies across a whole range of different industries.

Design is now a fundamental part of almost every aspect of how people live work and breath. Everything is designed from a toothbrush to every aspect of a web site. Compare visual design to architectural design—in both cases anybody can draw quite complex pictures. The resulting pictures could have stimulating and well drawn graphic elements. In both cases, the question is why does the world need designers? Exploring this question more deeply one can ask—does it make such a difference to how one perceives and understands a design when it is made by a professional rather than an amateur?

The trend in business intelligence is to design tools to provide flexibility and leave the world of visual design to the amateurs. Stephen Few comments in Information Dashboard Design⁵ that “Without a doubt I owe the greatest debt of gratitude to the many software vendors who have done so much to make this book necessary by failing to address or even contemplate the visual design needs of dashboards. Their kind disregard for visual design has given me focus, ignited my passion, and guaranteed my livelihood for years to come.”

Visual Designs within the described framework are well thought through in how the data is displayed. The described system allows good information visualization design concepts to be captured and delivered back to users as Visual Documents using unique data processing and analysis techniques.

Visual Designs

Method or Visual Design Classifications

According to this embodiment, ten Visual Design types are defined and incorporated into the described system. It will be understood that additional Visual Designs may be further defined including the creation of certain examples and actual Visual Designs for specific industry applications.

The visual design types include:

- Hierarchical
- Temporal
- Spatial
- Textual
- Virtual
- Structural
- Classical
- Pivotal
- Navigational
- Interactive

The following describes a method for the assessment of Visual Design quality. In assessing the quality of a Visual Design the following factors should be considered:

Alternative approaches—To assess the capability of a Visual Design it is important to contrast it with other visualization methods. In particular one should compare the visual design to a classical graph or table of numbers. This comparison is important as many data visualizations add considerable graphic weight but little informational value.

Visual simplicity—Looking at a visualization should not overload the mind. The simplicity of the visualization is important as it enhances interpretation and allows common understanding without training. Some visualizations require considerable training to be applied. In general, the described solution will not use these visual designs.

Data density—the density of data in a visualization is a critical measure of its overall value. Higher density visu-

alizations, if successful in maintaining their simplicity, have considerable potential to increase the flow of information to end users.

Volume of ink used—Is the visual design using negative space to show key information? This use of negative space allows lower volumes of ink to be used while showing the same or higher density of information. In addition, ink required is generally reduced as the number of “views” or pages of data is reduced to convey the same volume of data.

Capability to be illuminated with detail—In the end, data visualization becomes information visualization when the specific details are shown. The ability of a visualization to hold detailed information in specific places, often achieved with labels, is a key element in determining its value as an information visualization.

Visual Design Layers

There are seven defined Visual Design Layers which are set out diagrammatically as shown in FIG. 7. Other visual design layers may be added as appropriate.

These seven Visual Design Layers are described in the following table:

Visual Design Layer Type	Description
1. Embellishment Layers	Embellishment Layers have labels, symbology and/or other detailed information that is used to illuminate information that is displayed in the lower layers. The overlay can also include controls such as progress bars or spark-lines.
2. Selectable Layers	Selectable Layers are interactive and consist of items that can have associated data. On a retail spatial map it includes store locations as they have associated data. Selectable Layers are typically not obscured by thematic treatments.
3. Thematic Layers	Thematic Layers overlay colors or heatmaps on Special Layers. These thematic treatments become the core visual impact of the final Visual Document.
4. Transparent Thematic Layers	Transparent Thematic Layers are very similar to Thematic Layers (in fact are an alternative). The only difference is that they are graphically merged using a transparent overlay. For example, this kind of layer is necessary to overlay heatmaps on maps.google.com.
5. Special Layers	Special Layers construct the structure of the data. Specifically the Special Layer understands how to automatically draw the data so that other thematic treatments can be applied. Special Layers include mundane layers such as layers of polygons.
6. Context Layers	These are the lowest level of the visualization; they include background maps and other contextual information.
7. Context Map Layers	This is a type of context layer that is rendered from a map such as Google TM Maps, Yahoo TM Maps etc. This may be a road map, satellite map or any other map. It is past as a set of tiled images and as such can only be used as a Context Layer. Typically, a Transparent Thematic Layer will be used to display thematic data on a context map layer.

In terms of the Special Layer, two examples of Special Layers are set out below:

A. Classic Example of Special Layer: Voronoi Diagram
Source: Wikipedia⁶⁷

In mathematics, a Voronoi diagram, named after Georgy Voronoi, also called a Voronoi tessellation, a Voronoi decomposition, or a Dirichlet tessellation (after Lejeune Dirichlet), is a special kind of decomposition of a metric space determined by distances to a specified discrete set of objects in the space, e.g., by a discrete set of points.

In the simplest and most common case, in the plane, a given set of points S , and the Voronoi diagram for S is the partition of the plane which associates a region $V(p)$ with each point p from S in such a way that all points in $V(p)$ are closer to p than to any other point in S .

A Voronoi diagram can thus be defined as a Special Layer, where a set of polygons are generated from a set of points. The resulting polygon layer can then be subjected to thematic treatments, such as coloring.

B. Non Traditional Example of a Special Layer: Calendar

A calendar can be generated as a Special Layer for display of a temporal visual document. This Special Layer would require a 'start date' and an 'end date', most other information regarding the nature and structure of the Calendar could be determined automatically. The thematic layers would then use the structure of the calendar as a basis for thematic treatments such as coloring and contouring.

In an example from ENTROPÍA^{vii} a calendar is shown that can be created into a spiral. The structure and layout of this spiral will be the subject of considerable design discussions by information designers focused on issues such as aesthetics and clarity of information. The result of this discussion is a visual design of a spiral calendar Special Layer. This Special Layer can then be used for thematic treatments such as coloring.

As discussed in the above mentioned provisional patent application, an improved method of graphically representing data sets using tree methodologies has been developed as follows.

One purpose of a hierarchical visual design is to present large scale hierarchical data in one display. It is a picture for understanding, monitoring, exploring and analyzing hierarchical data.

Key elements of hierarchical visual designs are:

Data is hierarchical.

Structure of data can determine hierarchy.

They can be overlaid with connections.

This type of visualization may be automatically generated from a table of contents. This automatically generated hierarchy then becomes a special layer over which specific information can be overlaid.

A tree structure is one example of a hierarchical visual design. A typical hand drawn example of a tree structure is available at Wikipedia at http://en.wikipedia.org/wiki/Tree_structure. A hand drawn example of a tree structure of a French encyclopedia is shown. In this case the detail of the actual topics in the encyclopedia becomes the detail that allows the visual design to be useful.

It would be possible to overlay information on this visual design such as the average cost per page of building the document to show people the relationship between the natural structure of the document and some financial measure.

The Hierarchical Visual Design is a hierarchical diagram such as an organizational chart or a correlation matrix.

This Visual Design has at least one natural centre and typically has a higher density toward the fringes of the visualization. The Hierarchical Visual Design can typically be considered as a 'tree' structure. The nodes and vertices within the tree structure are best if they are generated automatically from a dataset. This tree structure is a good example of a Special Layer.

The development process includes building a tree that is optimized for this type of Visual Design including heat mapping techniques.

Large scale hierarchical data is represented using various techniques such as mapping to icons, shapes, colors and heights.

Typical uses include mapping of web pages, organizational charts, decision trees and menu options.

Generating a hierarchy from data may be completed with human intervention. For example, when looking at the hierarchy of an encyclopedia it would be possible to rearrange the density of information so there is no overlap between trees. The decision on how much overlap removal may be assisted with human input.

A hierarchy generated automatically often has more than one solution, for example an organizational chart can be meaningfully shown in a number of ways (e.g. horizontally or vertically). One key challenge in automating the tree production is to also describe the cartographic and layout decisions that can be captured into the Special Layers that build the tree.

A specific example is described below in relation to the determining of a hierarchy of data points in a textual type visual design.

Typically tree views have relationships between nodes. These relationships can be directional, bi-directional or non-directional. Relationships may also exist between multiple levels that complicate the process of drawing a useful visual document. Tree views may also be drawn without relationships between nodes. For example, in a fully connected network a relationship is not useful; however, the fully connected nodes may have a hierarchy that can be used to structure the Visual Document.

Potential Layouts include linked hierarchy, single value hierarchy and multiple value hierarchy.

The R-Tree is a traditional method for building an index of a spatially arranged dataset—this is also a natural Special Layer. By overlaying the creation of an R-Tree index over a data set, such as a self organizing map (SOM) or other dimension reducing representation, a hierarchy of the data can be constructed.

An example of this hierarchy is shown below in FIG. 8.

The generated R-Tree functions both as an index to navigation of the unstructured data and a summarization of the key components of the data.

It has not previously been known to use the dimensional indexing methodologies, such as the R-tree indexing, the Quad Tree or the Priority of R-tree as a display mechanism on its own. These indexing methods can be used to more easily enable the interpretation of visual representations, such as, for example, the interpretation of self organizing maps or other dimension reducing representation. Dimensional indexing can enable a user to interact with visualizations in a more meaningful manner by selecting data that is arranged in a hierarchical manner in the index and seeing where that data is located within the visualization. The combination of an indexing method and a display mechanism provide for a computationally efficient and technically advanced method for interacting with very large datasets.

Description of Dimensional Index: R-Tree

As defined in Wikipedia, R-trees are tree data structures that are similar to B-trees, but are used for spatial access methods i.e., for indexing multi-dimensional information; for example, the (X,Y) coordinates of geographical data. A common real-world usage for an R-tree might be: "Find all museums within 2 miles of my current location". The data structure splits space with hierarchically nested, and possibly overlapping, minimum bounding rectangles (MBR, otherwise known as bounding boxes, i.e. "rectangle", what the "R" in R-tree stands for).

Each node of an R-tree has a variable number of entries (up to some pre-defined maximum). Each entry within a

non-leaf node stores two pieces of data: a way of identifying a child node, and the bounding box of all entries within this child node.

The insertion and deletion algorithms use the bounding boxes from the nodes to ensure that “nearby” elements are placed in the same leaf node (in particular, a new element will go into the leaf node that requires the least enlargement in its bounding box). Each entry within a leaf node stores two pieces of information; a way of identifying the actual data element (which, alternatively, may be placed directly in the node), and the bounding box of the data element.

Similarly, the searching algorithms (for example; intersection, containment, nearest) use the bounding boxes to decide whether or not to search inside a child node. In this way, most of the nodes in the tree are never “touched” during a search. Like B-trees, this makes R-trees suitable for databases, where nodes can be paged to memory when needed.

A quad free is a tree whose nodes either are leaves or have 4 children. The children are ordered 1, 2, 3, 4.

The advantage of indexes, such as the R-Tree, Priority R-Tree or Quad-Tree indexes, is that they can classify a two (or higher) dimensional space into a hierarchical tree. This hierarchy is typically used to lookup values, for example to determine all the customers that are near to a store location.

Within the R-Tree, nodes at the same level of the tree may be balanced. This balancing provides for an insightful classification of data within the tree since each layer of the classification contains all data and an automatically generated hierarchy. A similar balancing is desirable for other kinds of spatial indexing functions such as the Priority tree.

Embodiments of the present invention are described herein with reference to a system adapted or arranged to perform a method for creating a visual representation of data points from metric data and determining a hierarchy of the data points.

In summary, the system includes at least a processor, one or more memory devices or an interface for connection to one or more memory devices, input and output interfaces for connection to external devices in order to enable the system to receive and operate upon instructions from one or more users or external systems, a data bus for internal and external communications between the various components, and a suitable power supply. Further, the system may include one or more communication devices (wired or wireless) for communicating with external and internal devices, and one or more input/output devices, such as a display, pointing device, keyboard or printing device.

The processor is arranged to perform the steps of a program stored as program instructions within the memory device. The program instructions enable the various methods of performing the invention as described herein to be performed. The program instructions may be developed or implemented using any suitable software programming language and toolkit, such as, for example, a C-based language. Further, the program instructions may be stored in any suitable manner such that they can be transferred to the memory device or read by the processor, such as, for example, being stored on a computer readable medium. The computer readable medium may be any suitable medium, such as, for example, solid state memory, magnetic tape, a compact disc (CD-ROM or CD-R/W), memory card, flash memory, optical disc, magnetic disc or any other suitable computer readable medium.

The system is arranged to be in communication with external data storage systems or devices in order to retrieve the relevant data.

It will be understood that the system herein described includes one or more elements that are arranged to perform the various functions and methods as described herein. The following portion of the description is aimed at providing the reader with an example of a conceptual view of how various modules and/or engines that make up the elements of the system may be interconnected to enable the functions to be implemented. Further, the following portion of the description explains in system related detail how the steps of the herein described method may be performed. The conceptual diagrams are provided to indicate to the reader how the various data elements are processed at different stages by the various different modules and/or engines.

It will be understood that the arrangement and construction of the modules or engines may be adapted accordingly depending on system and user requirements so that various functions may be performed by different modules or engines to those described herein.

It will be understood that the modules and/or engines described may be implemented and provided with instructions using any suitable form of technology. For example, the modules or engines may be implemented or created using any suitable software code written in any suitable language, where the code is then compiled to produce an executable program that may be run on any suitable computing system. Alternatively, or in conjunction with the executable program, the modules or engines may be implemented using any suitable mixture of hardware, firmware and software. For example, portions of the modules may be implemented using an application specific integrated circuit (ASIC), a system-on-a-chip (SoC), field programmable gate arrays (FPGA) or any other suitable adaptable or programmable processing device.

The methods described herein may be implemented using a general purpose computing system specifically programmed to perform the described steps. Alternatively, the methods described herein may be implemented using a specific computer system such as a data visualization computer, a database query computer, a graphical analysis computer, a gaming data analysis computer, a manufacturing data analysis computer, a business intelligence computer etc., where the computer has been specifically adapted to perform the described steps on specific data captured from an environment associated with a particular field.

According to this embodiment there is shown in FIG. 9 a conceptual system diagram where a data visualization system **901** includes various modules to provide visual designs (visual representations) using any number of methods as described herein. The visual designs are used to visually represent data in a form that is useful to the end user.

The system further includes a data retrieval module **903** which is arranged to retrieve metric data from a data storage module **909**, which is in communication with the data visualization system **901**.

The data provided as an input to the system may be of any suitable type of data, for example, real world data including, but not limited to, gaming or gambling data associated with a gaming environment such as a casino, event data, test or quality control data obtained from a manufacturing environment, business data retrieved from an accounting system, sales data retrieved from a company database, etc. All this data may be received by the system in real time in a cache memory or may be stored in a more permanent manner.

The data retrieval module is configured to enable the retrieval of data from a data storage module, which is in communication with the data visualization computing system. The data storage module may be any suitable type of data

storage system. For example, it may be an enterprise data warehouse (EDW), a data mart, a database, a storage array or any other suitable device or groups of devices that can store data for later retrieval. Further, the data storage module may be a cache memory used to temporarily store incoming data captured in real time.

A determination module 905 is provided that is arranged to determine data points for the visualization based on the metric data retrieved from the data storage module 909. That is, specific values associated with the retrieved metric data determine the data points that are to be displayed in the visualization. For example, the retrieved metric data may define specific manufacturing error or tolerance information associated with specific process levels used to manufacture a particular number of different products. Data values associated with the metric data may be obtained by the determination module and used to determine where the associated data points are to be rendered in the visualization.

A display module 907 arranges the determined data points to enable them to be displayed or rendered on a display device 911 according to a predetermined visual representation. The visual representation may be in any suitable format.

Further, the display module 907 arranges the data points into a plurality of meta groups in a hierarchical manner, and also arranges the meta groups into a plurality of layers, where each layer represents the data points at different levels of granularity. The hierarchy is determined from the metric data. Each layer in the hierarchy is calculated to "contain" or encompass the elements of the lower level. The hierarchy may be validly calculated such that the parent either has exclusive or non exclusive parenthood.

The system also includes a statistical distance determination module 917 that is controlled to determine a statistical distance between items within the meta groups by analyzing the metric data associated with the meta groups. Based on this determined statistical distance output from the statistical distance determination module 917, the display module 907 can arrange the items within the meta groups such that the most statistically similar items are arranged next to each other. This enables more useful information associated with the meta groups and metric data to be conveyed, which may not have otherwise been available to the user.

The statistical determination by the statistical distance determination module may be based on any suitable statistical calculation, such as, for example:

- i) analysis using a hierarchical force based algorithm,
- ii) determining the statistical distance between data groups,
- iii) determining the statistical distance based on the similarity of the data groups,
- iv) determining the statistical distance based on the sum of the squares of the data groups,
- v) determining the statistical distance based on the output of a heuristic algorithm, such as back propagation neural network values or genetic algorithm associations,
- vi) determining the statistical distance based on the output of a neural network, and
- vii) determining the statistical distance based on a correlation factor between the data groups, for example by calculating the value normalized by the sum of the squares.

Therefore, the display module may arrange the meta groups into at least two, possibly more, layers wherein the arrangement of meta groups in each layer represents all the data points at that level of granularity. That is, for that level of hierarchy, all the relevant data points associated with the metric data for that level are represented in the visualization.

The display module may also arrange the data points to be evenly distributed among all groups within each of the indi-

vidual layers, so that each meta group has an approximately equal number of data points as other meta groups in the same layer.

Optionally, the display module may organize a meta group in the first layer so that it represents a similar number of data points as a further meta group in the first layer. Also, the display module may arrange the data points so that common data points are associated with multiple groups within the same layer.

As an alternative to, or in conjunction with, the display module, further output modules may be provided to output the visualization. That is, the raw data retrieved by the data retrieval module is analyzed and arranged to provide the visualization output of the data in a specific format. The output data is provided to the display and/or further output modules to enable a user to visualize the raw data in a manner that conveys more useful or hidden information that would otherwise be lost, as well as to provide the user with the ability to correlate the data hierarchy with the visualization.

The further output module may be a printing device in communication with the described system to receive print control data so that representations of the data may be printed on any suitable print medium. Alternatively, the further output module may be an interface that enables the data output to be interfaced with other data handling modules or storage devices.

The system further includes an activation detection module 913 that is arranged to detect the activation of a meta group. The activation detection module detects an end user selecting a meta group on the visualization through the use of an input device 915. The input device in this embodiment is a pointing device, such as a computer mouse, that enables the system to recognize when the user has selected a particular area on the visual representation that corresponds with a meta group.

Upon the activation detection module detecting activation of the meta group, the display module is controlled to provide an indication of the data points associated with the activated meta group in the visual representation. That is, the system determines which meta group has been selected in the visual representation and determines from data within the data storage module which portions of the visual design (representation) are associated with the selected meta group. The relevant portion of the visual design is modified by the rendering engine to indicate to the user which meta groups were selected. For example, the relevant portion of the visual design may be modified so that it flashes, or changes color, etc.

The display module may also overlay a graphical representation of the hierarchy over the visual representation, or display it side by side with the visual representation. According to this embodiment, the display module arranges the groups as a graphical tree representation along with the visual representation.

Further, the system may be arranged to determine the hierarchy of the metric data at substantially the same time as the visual representation is being created by the visualization system. This is accomplished by calculating the distribution of nodes in the representation as they are rendered. This has the advantage of "just in time" calculation and may result in a reduced calculation overhead. For example, some nodes may not require adjustment as they are not visible in the visualization.

Therefore, the display module may arrange icons to be displayed in a tree structure within the visual representation, where the icons represent the meta groups. Each meta group may be arranged for display by the display module using a predefined icon where a relative difference between each icon

indicates a relative difference in the metric data associated with each meta group. For example, the relative size or shape of an icon can indicate that the data associated with each meta group is or a larger or smaller value.

As an example, the icon generated by the display module for display on the display device may be a geometric shape, such as a triangle, rectangle or quadrangle. It will be understood that, as an alternative, the display module may create other shapes to represent the data.

The system may also be arranged to enable the display module to display overlapping icons. The overlapping icons may be used to represent a similarity in metric data associated with each meta group.

It will be understood that the data points may represent any type of information that can be visually represented in a hierarchy. For example, product hierarchies, customer segmentation, document table of contents etc. may be represented using the data points.

The display module is further arranged to adapt the visual representation so that it represents data associated with a physical world entity in the form of the physical world entity itself. For example, the physical world entity may be one of a geographic location, road location or road intersection, where the physical arrangement of the locations and intersections are represented in the visualization. Further, the geographic location may be one or more stores or distribution centers, where the representation shows the physical or geographic location of those stores or centers. In order to display the physical location of these elements, the system may incorporate various other databases and systems to retrieve geographical location information and map data in order to render the visualization.

The display module may utilize spatial algorithms to determine the best way to display the data hierarchy on the visualization. That is, the visual representation may be adapted by the system in order to reduce the amount of display space required to display the hierarchy.

According to the embodiments described above, a method is provided that enables an end user to: create a visual representation of data points from metric data, and determine a hierarchy of the data points for displaying with the visual representation.

As shown in FIG. 10, this is carried out by retrieving the metric data in step 1001, determining the data points based on the retrieved metric data in step 1003, arranging the data points to be displayed according to a predetermined visual representation in step 1005, arranging the data points into a plurality of meta groups in a hierarchical manner in step 1007, and arranging the meta groups into a plurality of layers in step 1009, where each layer represents the data points at different levels of granularity.

Further, the method includes a statistical distance determination module that is controlled to determine a statistical distance between items in the meta groups in step 1011, and the display module arranging the items within the meta groups based on the determined statistical distance in step 1013.

In addition, the method includes the meta groups being arranged into two or more layers, wherein the groups in each layer represent all the data points at that level of granularity.

Also, data points may be arranged to be evenly distributed among all groups within a single layer. Further, a group in the first layer may represent a similar number of data points as a further group in the first layer.

Common data points may be associated with multiple groups within the same layer. Thus, when a user activates one of the groups, for example by clicking on the group in the representation, an indication is made in the visual representation of the data points associated with the activated group.

The hierarchy or granularity may be created from the metric data at substantially the same time as the visual representation is created.

A graphical representation of the hierarchy may be displayed over the visual representation. The groups may be arranged as a graphical tree representation.

Referring to FIG. 8, an example of a visualization in the form of a tree diagram 800 is shown. The tree diagram visualization is based on R-Tree principles, and may be used by the user to navigate around a Visual Document 825, as well as to proportionately display the significance of the leafs below each hierarchical layer. Data points being visualized are arranged in a plurality of meta groups in a hierarchy. The meta groups are represented by icons, for example, rectangles.

Size is used when drawing the icons in order to represent each level in a manner that indicates the size or significance of the levels below. For example, at level 1 805, the width of the icon (a rectangle in this example) is calculated to be equal to the sum of the widths of the rectangles associated with it in its child level 810. The widths of the rectangles in level 810 are equal to the sum of the widths of the associated rectangles in its child level 815. The width of a single rectangle (R1.1.2) with children in level 815 is equal to the widths of its children (R1.1.2.1 and R1.1.2.2) in level 820. However the widths of each rectangle in each level vary depending on the associated data they represent relative to the other nodes in their level and the parent node.

Although in this example, only the width of each rectangle in the tree representation has been varied to represent the relative relationship between the different rectangles (or meta groups) in each level, it will be understood that other methods may also be employed to display the relative relationship. For example, the relative overall size of each rectangle may be used to indicate the relative number of items associated with that node. Alternatively, the height of each icon may be varied depending on the values associated with the node.

As a further alternative, other shapes, lines, colors or other graphical methods may be used to represent each icon associated with each node in the tree in order to indicate the relative difference between nodes, i.e. between each meta group. For example, each meta group may be represented using an icon of any suitable geometric shape, such as a triangle, rectangle or quadrangle, for example, or any other geometric shape of any number of sides, or of an irregular shape. The meta groups may also be represented using organic shaped icons.

It will be understood that one or a combination of the above described methodologies could be implemented at any one time to graphically depict the metric data in different ways simultaneously. Further, each of the nodes may be arranged to overlap with other nodes to indicate that there is similar or overlapping data set associated with the nodes.

Symbols, shading or shapes within rectangles may also be used to communicate information as illustrated in node R1.2.2 835, and R1.1.1 845. For example, if the data being represented has been detected to exceed a preset threshold by the system, the area may be indicated.

The position of the node may also used to communicate additional information. For example, a region of high importance 840 indicates that nodes in that area are significant.

If a node on the R-Tree is selected by an end user, the associated area, such as an individual Visual Design 830 within the associated Visual Document 825 is highlighted. Thus a user can easily determine from the hierarchically arranged metric data where the associated data points are located within the visualization, as well as determining the relative importance of each group of metric data.

Any changes that are made to the Visual Document 825 are reflected in the Tree Diagram 800 and vice versa. This update

facility ensures that the tree hierarchy is kept in synchronization with the displayed visualization.

Therefore, the data visualization techniques described herein transform the raw data from the data storage module into both a hierarchical tree like representation conveying structural or hierarchical information about the data along with a visual representation conveying value information obtained from the data. This therefore enables further or hidden information within the raw data to be visually represented in a manner that conveys the information to a user in an efficient manner.

It will be understood that the graphical representation being produced may be adapted specifically so that reproductions of the representation minimize the use of production printing materials. That is, by producing a reproduction or representation that has one or more characteristics to summarize a complex problem a number of technical advantages are immediately provided. For example, the characteristics may include the limitation of the size of the representation, the use of a minimum amount of ink, or the creation of the representation using a minimal or bounded area space or minimum amount of time. These characteristics then may solve the problem of excessive consumption of consumable items, by reducing the required consumption of consumables such as paper and ink resources, as well as reducing the energy required to produce the printouts due to the ability to provide the required information on media space of a smaller size.

According to one example, the representation may be recreated by efficiently utilizing white space rather than producing the representation merely by applying ink to create the representation. Also the representation may be arranged by the display module of the system to reduce the amount of display space required to display the hierarchical information.

As a system option, each of the nodes of the index may be linked with summarized or cached query results. These stored values may be used by the system to rapidly calculate and display business performance drivers (BPDs) as the user interacts with the tree by detecting which nodes are selected by the user and retrieving the associated information.

FIG. 11 shows an example of how the herein described system may be incorporated within a gaming environment. The gaming environment consists of a number of gaming machines **1101** and electronic tables **1103** (among other electronic gaming devices) that are adapted to communicate electronically with other systems using any suitable protocols, such as data packet protocols.

The gaming environment further includes a number of electronic cashier devices **1105** and ATMs **1107** which are in communication via a Wide Area Network **1109** with one or more financial databases **1111**.

Data from the gaming machines **1101** and electronic tables **1103** are transferred to a reward program database **1113** and customer database **1115**. It will be understood that these two databases may be combined into a single database.

Data from the cashier devices are also transferred to the reward program database **1113** and customer database **1115**. The databases **1113** and **1115** are in communication with a central hotel management system **1117** that oversees the operation of the gaming environment, including the activities of customers in other areas of a casino, such as shops, hotels, spas etc.

The system **1119** described herein is in communication with the reward program database **1113**, customer database **1115** and central hotel management system **1117** so the system can retrieve all necessary data about the activities within

the gaming environment. The various embodiments as described herein are employed by the system **1119** to provide an output **1121**.

Glossary

Term	Definition
Agile Development	Agile software development is a conceptual framework for software engineering that promotes development iterations throughout the life-cycle of the project. There are many agile development methods; most minimize risk by developing software in short amounts of time. Software developed during one unit of time is referred to as an iteration, which may last from one to four weeks. Each iteration is an entire software project: including planning, requirements analysis, design, coding, testing, and documentation. An iteration may not add enough functionality to warrant releasing the product to market but the goal is to have an available release (without bugs) at the end of each iteration. At the end of each iteration, the team re-evaluates project priorities. Wikipedia ^{xiii}
BPD Packages	A BPD Package is made up from a set of related BPDs. This relationship (between a BPD Package and its BPDs) is defined using metadata. BPD Packages can be thought of as the Visual Document's vocabulary.
Catalog	The described catalog is used to store permanent and temporary objects that are necessary for creation and storage of Visual Documents. These may include Visual Designs, BPDs, Configuration tools and other objects. There may be multiple catalogs of different types (such as - database, flat file) which are configured by an integrator - dependent on customer requirements. All items in a catalog are identified by a unique ID and can only be accessed by those with the correct authorization.
Data Packages	Data Packages contain data that can be sold with subscription or service provision including an associated managed dataset. For example, census data will be available as a Data Package; this Data Package will enable the described solution users to interact and use a slowly changing dataset called census. (Census data can be updated after each census and is often modeled between each census).
Dimension	Dimensional modeling always uses the concepts of facts (sometimes referred to as measures) and dimensions. Facts are typically (but not always) numeric values that can be aggregated, and dimensions are groups of hierarchies and descriptors that define the facts. For example, sales amount is a fact; timestamp, product, register#, store#, etc. are elements of dimensions. Wikipedia ^x
Dimensional Modeling	DM is a logical design technique that seeks to present the data in a standard, intuitive framework that allows for high-performance access. It is inherently dimensional, and it adheres to a discipline that uses the relational model with some important restrictions. Every dimensional model is composed of one table with a multipart key, called the fact (sometimes referred to as measures) table, and a set of smaller tables called dimension tables. Each dimension table has a single-part primary key that corresponds exactly to one of the components of the multipart key in the fact table. Intelligent Enterprise ⁷
Enterprise Java Beans (EJBs)	In a typical J2EE application, Enterprise JavaBeans (EJBs) contain the application's business logic and live business data. Although it is possible to use standard Java objects to contain your business logic and business data, using EJBs addresses many of the issues you would find by using simple Java objects, such as scalability, lifecycle management, and state management. Wikipedia ^{xii}
Fact	Dimensional modeling always uses the concepts of facts (sometimes referred to as measures) and dimensions. Facts are typically (but not always) numeric values that can be aggregated, and dimensions are groups of hierarchies and descriptors that define the facts. For example, sales amount is a fact; timestamp, product, register#, store#, etc. are elements of dimensions. Wikipedia ^{xii}

Term	Definition
IIOP (Internet Inter-ORB Protocol)	IIOP (Internet Inter-ORB Protocol) is a protocol that makes it possible for distributed programs written in different programming languages to communicate over the Internet. SearchCIO - Midmarket ^{xiii}
KML	Keyhole Markup Language. Google TM Earth is a geographic browser -- a powerful tool for viewing, creating and sharing interactive files containing highly visual location-specific information. These files are called KMLs (for Keyhole Markup Language): what HTML is to regular Internet browsers, KML is to geographic browsers. You can open KML files in both Google TM Earth and Google TM Maps, as well as in many other geographic browsers. Google TM Maps ^{xiv}
MDT	The average time that a system is non-operational. This includes all time associated with repair, corrective and preventive maintenance; self imposed downtime, and any logistics or administrative delays. The difference between MDT and MTTR (mean time to repair) is that MDT includes any and all delays involved; MTTR looks solely at repair time. Wikipedia ^{xv}
Metadata	Metadata describes how data is queried, filtered, analyzed, and displayed in the described solution. In general terms, metadata is data about data. For example, in a library the metadata (pertaining to the catalog of books) could be - the title of the book, the author(s), categories (e.g. reference, fiction, non-fiction etc), physical location. This metadata can be used in searches, directories etc to help users locate books.
MTBF	Mean time between failures (MTBF) is the mean (average) time between failures of a system, and is often attributed to the 'useful life' of the device i.e. not including 'infant mortality' or 'end of life'. Calculations of MTBF assume that a system is 'renewed', i.e. fixed, after each failure, and then returned to service immediately after failure. The average time between failing and being returned to service is termed mean down time (MDT) or mean time to repair (MTTR). MTBF = (downtime - uptime)/number of failures. Wikipedia ^{xvi}
MTTR	Mean Time to Recovery - the average time that a device will take to recover from a non-terminal failure. Examples of such devices range from self-resetting fuses (where the MTTR would be very short, probably seconds), up to whole systems which have to be replaced. The MTTR would usually be part of a maintenance contract, where the user would pay more for a system whose MTTR was 24 hours, than for one of, say, 7 days. This does not mean the supplier is guaranteeing to have the system up and running again within 24 hours (or 7 days) of being notified of the failure. It does mean the average repair time will tend towards 24 hours (or 7 days). A more useful maintenance contract measure is the maximum time to recovery which can be easily measured and the supplier held accountable. Wikipedia ^{xvii}
OLAP	On Line Analytical Processing. OLAP performs multidimensional analysis of business data and provides the capability for complex calculations, trend analysis, and sophisticated data modeling. OLAP enables end-users to perform ad hoc analysis of data in multiple dimensions, thereby providing the insight and understanding the need for better decision making. Paris TM Technologies ^{xviii}
Planogram	A planogram is a diagram of fixtures and products that illustrates how and where retail products should be displayed, usually on a store shelf in order to increase customer purchases. They may also be referred to as planograms, plan-o-grams, schematics (archaic) or POGs. A planogram is often received before a product reaches a store, and is useful when a retailer wants multiple store displays to have the same look and feel. Often a consumer packaged goods manufacturer will release a new suggested planogram with their new product, to show how it relates to existing products in said category. Planograms are used nowadays in all kind of retail areas. A planogram defines which product is placed in which area of a shelving unit and with which quantity. The rules

Term	Definition
5 Request Queue	and theories for the creation of a planogram are set under the term of merchandising. Wikipedia ^{xx} The Request Queue manages Visual Documents requests generated by a user or the scheduler. As requests are processed, the Visual Document maintains various statuses until the Visual Document is complete and available to be viewed by a user.
10 SaaS	Software as a Service. A software application delivery model where a software vendor develops a web-native software application and hosts and operates (either independently or through a third-party) the application for use by its customers over the Internet. Customers do not pay for owning the software itself but rather for using it. Wikipedia ^{xx}
15 Scrum	Scrum is an agile process that can be used to manage and control software development. With Scrum, projects progress via a series of iterations called sprints. These iterations could be as short as 1 week or as long as 1 month. Scrum is ideally suited for projects with rapidly changing or highly emergent requirements. The work to be done on a Scrum project is listed in the Product Backlog, which is a list of all desired changes to the product. At the start of each sprint a Sprint Planning Meeting is held during which the Product Owner prioritizes the Product Backlog and the Scrum Team selects the tasks they can complete during the coming Sprint. These tasks are then moved from the Product Backlog to the Sprint Backlog. Each day during the sprint a brief daily meeting is held called the Daily Scrum, which helps the team stay on track.
20 Self Organizing Maps (SOM)	At the end of each sprint the team demonstrates the completed functionality at a Sprint Review Meeting. A type of artificial neural network that is trained using unsupervised learning to produce a low-dimensional (typically two dimensional), representation of the input space of the training samples, called a map. The map seeks to preserve the topological properties of the input space. Wikipedia ^{xxi}
25 Servlets	Servlets are modules of Java code that run in a server application (hence the name "Servlets", similar to "Applets" on the client side) to answer client requests. Servlets are not tied to a specific client-server protocol but they are most commonly used with HTTP and the word "Servlet" is often used in the meaning of "HTTP Servlet". Servlets make use of the Java standard extension classes. Since Servlets are written in the highly portable Java language and follow a standard framework, they provide a means to create sophisticated server extensions in a server and operating system independent way. Typical uses for HTTP Servlets include: 1. Processing and/or storing data submitted by an HTML form. 2. Providing dynamic content, e.g. returning the results of a database query to the client. 3. Managing state information on top of the stateless HTTP, e.g. for an online shopping cart system which manages shopping carts for many concurrent customers and maps every request to the right customer.
30 Subject Matter Expert (SME)	Servlet Essentials ^{xxii} The Subject Matter Expert is that individual who exhibits the highest level of expertise in performing a specialized job, task, or skill within the organization. Six Sigma ^{xxiii}
35 WebSphere	WebSphere is an IBM TM brand of products that implement and extend Sun's JavaTwoEnterpriseEdition (J2EE) platform. The Java-based application and transaction infrastructure delivers high-volume transaction processing for e-business and provides enhanced capabilities for transaction management, as well as security, performance, availability, connectivity, and scalability. IBM TM WebSphere Product Pages ^{xxiv}
40 Subject Matter Expert (SME)	
45 WebSphere	
50 Subject Matter Expert (SME)	
55 WebSphere	
60 Subject Matter Expert (SME)	
65 WebSphere	

Further Embodiments

It will be understood that the embodiments of the present invention described herein are by way of example only, and that various changes and modifications may be made without departing from the scope of invention.

In an alternative embodiment, the display module is arranged to build the hierarchy representation using a quartering algorithm. The quartering algorithm operates to ensure that each node in each layer contains approximately one quarter of the data of the parent node. Alternatively, the quartering algorithm may operate to ensure each node contains objects whose cumulative values represent approximately one quarter of the value of the parent node. With some datasets this may be more efficient as it reduces the computational complexity because the calculation of the number of nodes at each level may not be required. This method may tend to produce deeper hierarchical structures as an output.

The values used by the quartering algorithm are preferably values that are calculated from a classifying value that is not the count of the number of items associated with the node. This “secondary” value may include revenue, house hold income or frequency, for example. Further, the distribution to each node may be calculated approximately using the quartering algorithm so that the child node contains approximately 25% of the value of the parent.

As a further alternative, the representation may be built using a tree balancing module and algorithm. For example, each node in each layer may be arranged by the tree balancing module to contain a defined subset of the data of the parent node. Alternatively, each node may be arranged by the tree balancing module to contain objects whose cumulative values represent a subset of the value of the parent node. Further, the values may be calculated from a classifying value that is not the count of the number of items. This “secondary” value may include revenue, house hold income or frequency, for example. Further, the distribution to each node may be calculated approximately, so that each sibling node contains approximately the same value.

According to a further embodiment, a Quad Tree may be utilized to provide for a non visually balanced navigational and visualization method.

In a further embodiment, a dimensional reduction technique, such as self organizing maps or simulated annealing for example, may be applied to the input metric data to generate a visual representation of a dimensionally reduced data set. A spatial index may then be visually generated by the system using, for example, an R-Tree classification. The spatial index visualization may then provide a mechanism for summarization and exploration annotation of each of the nodes. This exploration technique improves the efficiency in exploration of the dimensionally reduced image. This efficiency greatly increases the information density on the display enabling, for example, a printout of limited size to summarize a complex problem. This therefore solves the further problem of reducing the consumption of consumable items, such as paper and ink resources, as well as reducing the energy required to produce the printouts due to their smaller size.

In a further embodiment, the user may dynamically reallocate items between nodes in the spatial index, by picking and placing the items in a new location using a computer input device, such as a mouse for example. The system may detect the selection and position of the placement and make the associated changes in the data sets. This reallocation enables dynamic classification and summarization of specific items within the matrix in a more intuitive manner. The reallocation could be managed in a separate and balanced index tree.

ⁱ Wikipedia; “Heat map”; Date Accessed, Jun. 10, 2008; http://en.wikipedia.org/wiki/Heat_map

ⁱⁱ Nasdaq; Nasdaq-100 Dynamic Heatmap; Date Accessed, Jun. 10, 2008; http://screening.nasdaq.com/heatmaps/heatmap_100.asp

ⁱⁱⁱ Weather.com; Date Accessed, Jan. 31, 2008; <http://weather.com/>.

^{iv} Information Aesthetics; Information Aesthetics; “travel time maps”; Date Accessed, Jan. 31, 2008; <http://infosthetics.com/archives/locative/>

^v Few, Stephen—from white paper “BizViz: The Power of Visual Business Intelligence”—Mar. 7, 2006. www.perceptualedge.com

^{vi} Wikipedia; Wikipedia; “Voronoi Diagram”; Date Accessed, Jan. 31, 2008; http://en.wikipedia.org/wiki/Voronoi_diagram.

^{vii} ENTROPÍA; ENTROPÍA; “Más tiempo”; Date Accessed, Jan. 31, 2008; <http://www.luispabon.com/entropia/index.php?entry=entry071129-145959>.

^{viii} Wikipedia; Wikipedia; “Agile Software Development”; Date Accessed, Jan. 30, 2008; http://en.wikipedia.org/wiki/Agile_software_development

^{ix} Wikipedia; Wikipedia; “Dimensional Modeling”; Date Accessed: Apr. 10, 2008; http://en.wikipedia.org/wiki/Dimensional_modeling

^x Kimball, Ralph, A Dimensional Modeling Manifesto; Date Accessed: Apr. 10, 2008; <http://www.dbmsmag.com/9708d15.html>

^{xi} Sam’s Publishing; developer.com Gamelan™; “Introduction to EJB’s”; <http://www.developer.com/java/ejb/article.php/1434371>.

^{xii} Wikipedia; Wikipedia; “Dimensional Modeling”; Date Accessed: Apr. 10, 2008; http://en.wikipedia.org/wiki/Dimensional_modeling

^{xiii} Gilbert, Cheryl, et. al.; SearchCIO—Midmarket; “IIOP”; date Accessed, Jan. 30, 2008; http://searchcio-midmarket.techtarget.com/sDefinition/0,,sid183_gci214019,00.html.

^{xiv} Google; Google Maps; “KML Gallery: Explore the Earth on Google”; Date Accessed, Jan. 30, 2008; <http://earth.google.com/gallery/>

^{xv} Wikipedia; “Mean down time”; Date Accessed, Jan. 30, 2008; http://en.wikipedia.org/wiki/Mean_down_time.

^{xvi} Wikipedia; Wikipedia; “Mean time between failures”; Date Accessed, Jan. 30, 2008; http://en.wikipedia.org/wiki/Mean_time_between_failures.

^{xvii} Wikipedia; Wikipedia; “Mean time to recovery”; Date Accessed, Jan. 30, 2008; http://en.wikipedia.org/wiki/Mean_time_to_recovery.

^{xviii} Paris Technologies, Inc.; Paris Technologies; “OLAP”; Date Accessed, Jan. 30, 2008; <http://www.olap.com>.

^{xix} Wikipedia; Wikipedia; “Planogram”; Date Accessed, Jan. 30, 2008; <http://en.wikipedia.org/wiki/Planogram>.

^{xx} Wikipedia; Wikipedia; “Software as a Service”; Date Accessed, Jan. 30, 2008; http://en.wikipedia.org/wiki/Software_as_a_Service.

^{xxi} Wikipedia; Wikipedia; “Self-organizing map”; Date Accessed, Jan. 30, 2008; http://en.wikipedia.org/wiki/Self-organizing_map.

^{xxii} Zeiger, Stefan, Servlet Essentials, Version 1.3.6—Nov. 4, 1999; Date Accessed, Jan. 30, 2008

^{xxiii} Six Sigma; Subject Matter Expert—SME; Date Accessed: Jan. 30, 2008; http://www.isixsigma.com/dictionary/Subject_Matter_Expert_-_SME-396.htm

^{xxxiv} IBM; WebSphere Product Pages; "WebSphere software"; Date Accessed: Jan. 30, 2008; http://www-306.ibm.com/software/websphere/?pgel=ibmhzn&cm_re=masthead-_products-_sw-websphere.

What we claim is:

1. In a data visualization system, a method of creating a visual representation of data points from metric data, wherein the method includes the steps of:

a data retrieval module retrieving the metric data from a data storage module in communication with the data visualization system,

a determination module determining the data points based on the retrieved metric data,

a display module arranging the data points for displaying on a display device according to a predetermined visual representation,

the display module arranging the data points into a plurality of meta groups in a hierarchical manner,

the display module arranging the meta groups into a plurality of layers, where each layer represents the data points at different levels of granularity,

the display module arranging the data points to be evenly distributed among all meta groups within a single layer, a statistical distance determination module determining a statistical distance between data points in the meta groups, and

the display module arranging the data points within the meta groups based on the determined statistical distance, the most statistically similar data points being arranged next to each other.

2. The method of claim 1 further including the step of the display module arranging the meta groups into two or more layers, wherein the meta groups in each layer represent all the data points at that level of granularity.

3. The method of claim 1 wherein the meta groups are arranged into a plurality of layers where one or more child layers are the same size as their parent layers.

4. The method of claim 3, wherein common data points are associated with multiple meta groups within the same layer.

5. The method of claim 1, further including the step of an activation detection module detecting the activation of a meta group, whereupon detection of activation of the meta group, the display module indicates the data points associated with the activated meta group in the visual representation.

6. The method of claim 5 further including the step of the activation detection module detecting an end user selecting a meta group using an input device.

7. The method of claim 1, further including the step of the display module overlaying a graphical representation of the hierarchy over the visual representation.

8. The method of claim 7, further including the step of the display module arranging the meta groups as a graphical tree representation.

9. The method of claim 1, further including the step of determining the hierarchy from the metric data at substantially the same time as the visual representation is created.

10. The method of claim 1 further including the step of the display module arranging icons that represent the meta groups to be displayed in a tree structure within the visual representation.

11. The method of claim 10 further including the step of the display module displaying each meta group using a pre-defined icon where a relative difference between each icon indicates a relative difference in the metric data associated with each meta group.

12. The method of claim 11, wherein the icon is a geometric shape.

13. The method of claim 12, wherein the icon is one of a triangle, rectangle or quadrangle.

14. The method of claim 10 further including the step of the display module displaying overlapping icons to represent a similarity in metric data associated with each meta group.

15. The method of claim 1 including the use of a product hierarchy as the data points.

16. The method of claim 1 including the use of customer segmentation as the data points.

17. The method of claim 1 including the use of a document table of contents as the data points.

18. The method of claim 1 further including the step of adapting the representation so its reproduction minimizes the use of production printing materials.

19. The method of claim 1 further including the step of the display module adapting the representation so it represents a physical world entity.

20. The method of claim 19 wherein the physical world entity is one of a geographic location, road location or road intersection.

21. The method of claim 20 wherein the geographic location is the location of one or more stores or distribution centers.

22. The method of claim 1 further including the step of the display module adapting the representation to reduce the amount of display space required to display the hierarchy.

23. The method of claim 1 further including the step of the display module building the representation using a quartering algorithm where each node in each layer contains one quarter of the data of the parent node or each node contains objects whose cumulative values represent one quarter of the value of the parent node.

24. The method of claim 23 where the values are calculated from a classifying value that is not the count of the number of data points.

25. The method of claim 23 where the distribution to each node is calculated approximately, so that the child node contains about 25% of the value of the parent.

26. The method of claim 1 further including the step of building the representation using a tree balancing algorithm where each node in each layer contains a defined subset of the data of the parent node or each node contains objects whose cumulative values represent a subset of the value of the parent node.

27. The method of claim 26 where the values are calculated from a classifying value that is not the count of the number of data points.

28. The method of claim 26 where the distribution to each node is calculated approximately, so that each sibling node contains approximately the same value.

29. A data visualization system for creating a visual representation of data points from metric data, wherein the system includes:

a data retrieval module arranged to retrieve the metric data from a data storage module in communication with the data visualization system,

a determination module arranged to determine the data points based on the retrieved metric data, and

a display module executed by a processing device and arranged to:

arrange the data points for displaying on a display device according to a predetermined visual representation, arrange the data points into a plurality of meta groups in a hierarchical manner,

arrange the meta groups into a plurality of layers, where each layer represents the data points at different levels of granularity, and

41

arrange the data points to be evenly distributed among all meta groups within a single layer, and a statistical distance determination module arranged to determine a statistical distance between data points in the meta groups, wherein the display module is further arranged to arrange the data points within the meta groups based on the determined statistical distance.

30. The system of claim 29 wherein the display module is further arranged to arrange the meta groups into two or more layers, wherein the meta groups in each layer represent all the data points at that level of granularity.

31. The system of claim 29 further including an activation detection module arranged to detect the activation of a meta group, whereupon detection of activation of the meta group, the display module is further arranged to indicate the data points associated with the activated meta group in the visual representation.

42

32. The system of claim 29 wherein the display module is further arranged to overlay a graphical representation of the hierarchy over the visual representation.

33. The system of claim 29 wherein the display module is further arranged to arrange icons that represent the meta groups to be displayed in a tree structure within the visual representation.

34. The system of claim 29 wherein the display module is further arranged to build the representation using a quartering algorithm where each node in each layer contains one quarter of the data of the parent node or each node contains objects whose cumulative values represent one quarter of the value of the parent node.

35. The system of claim 29 wherein the meta groups are arranged into a plurality of layers where one or more child layers are the same size as their parent layers.

* * * * *